

Catching the k -NAESAT Threshold

Amin Coja-Oghlan^{1*} and Konstantinos Panagiotou²

¹ University of Warwick, Mathematics and Computer Science, Zeeman building, Coventry CV4 7AL, UK
a.coja-oghlan@warwick.ac.uk

² Max-Planck-Institute for Informatics, Campus E1.4, 66123 Saarbrücken, Germany
kpanagio@mpi-inf.mpg.de

Abstract. The best current estimates of the thresholds for the existence of solutions in random constraint satisfaction problems (‘CSPs’) mostly derive from the *first* and the *second moment method*. Yet apart from a very few exceptional cases these methods do not quite yield matching upper and lower bounds. According to deep but non-rigorous arguments from statistical mechanics, this discrepancy is due to a change in the geometry of the set of solutions called *condensation* that occurs shortly before the actual threshold for the existence of solutions (Krzakala, Montanari, Ricci-Tersenghi, Semerjian, Zdeborová: PNAS 2007). To cope with condensation, physicists have developed a sophisticated but non-rigorous formalism called *Survey Propagation* (Mézard, Parisi, Zecchina: Science 2002). This formalism yields precise conjectures on the threshold values of many random CSPs. Here we develop a new *Survey Propagation inspired second moment method* for the random k -NAESAT problem, which is one of the standard benchmark problems in the theory of random CSPs. This new technique allows us to overcome the barrier posed by condensation rigorously. We prove that the threshold for the existence of solutions in random k -NAESAT is $2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + \varepsilon_k$, where $|\varepsilon_k| \leq 2^{-(1-o_k(1))k}$, thereby verifying the statistical mechanics conjecture for this problem.

Key words: random structures, phase transitions, k -NAESAT, second moment method, Survey Propagation.

* Supported by EPSRC grant EP/G039070/2 and ERC Starting Grant 278857-PTCC (FP7).

1 Introduction

Over the past decade, physicists have developed sophisticated but non-rigorous techniques for the study of random constraint satisfaction problems (‘CSPs’) such as random k -SAT or random graph k -coloring [27,29]. This work has led to a remarkably detailed *conjectured* picture, according to which various phase transitions affect both the combinatorial and computational nature of random problems. By now, some of these predictions have been turned into rigorous theorems. Examples include results on the “shattering” of the solution space [1,7], work on (non-)reconstruction and sampling [18,24,30], and even new algorithms for random CSPs [9,19]. Many of these contributions have led to the development of new rigorous techniques. Indeed, it seems fair to say that, combined, these results have advanced our understanding of random CSPs quite significantly.

However, thus far substantial bits of the statistical mechanics picture have eluded all rigorous attempts. Perhaps most importantly, apart from a very few special cases, the precise thresholds for the existence of solutions in random CSPs have not been pinned down exactly. While rigorous upper and lower bounds can be derived via the *first* and the *second moment method* [5], these bounds do not quite match in most examples, including prominent ones such as random k -SAT or random graph k -coloring. In fact, the statistical mechanics techniques suggest a striking explanation for this discrepancy, namely the existence of a *condensation phase* shortly before the threshold for the existence of solutions. In this phase, a crucial necessary condition for the success of the (standard) second moment method is violated. Indeed, in statistical mechanics a deep formalism called *Survey Propagation* (‘SP’) has been developed expressly to deal with condensation. While SP is primarily an analysis technique, an off-spin has been the *SP guided decimation* algorithm, which seems highly successful at solving random CSPs experimentally.

In this paper we propose a new *SP-inspired second moment method* that allows us to overcome the barrier posed by condensation. The specific problem that we work with is random k -NAESAT, one of the standard benchmark problems in the theory of random CSPs. Random k -NAESAT is technically a bit simpler than random k -SAT due to a certain symmetry property, but computationally and structurally both problems have strong similarities. We determine the threshold for the existence of solutions in random k -NAESAT up to an additive error that tends to zero exponentially with k . This is the first time that the threshold in any random CSP of this type can be calculated with such accuracy. While from a technical viewpoint k -NAESAT is perhaps the simplest example of a random CSP that exhibits condensation, our proof technique rests on a rather generic approach. Therefore, we believe that with additional technical work our approach can be extended to many other problems, including random k -SAT or random graph k -coloring.

To define random k -NAESAT formally, let $k \geq 3$ and $n > 0$ be integers and let $V = \{x_1, \dots, x_n\}$ be a set of Boolean variables. For a fixed real $r > 0$ we let $m = m(n) = \lceil rn \rceil$. Further, let $\Phi = \Phi_k(n, m)$ be a propositional formula obtained by choosing m clauses of length k over V uniformly and independently at random among all $(2n)^k$ possible clauses. We say that an assignment $\sigma : V \rightarrow \{0, 1\}$ is an *NAE-solution* (a “solution”) if each clause has both a literal that evaluates to ‘true’ under σ and one that evaluates to ‘false’. In other words, both σ and its inverse $\bar{\sigma} : x_i \mapsto 1 - \sigma(x_i)$ are satisfying assignments of the Boolean formula Φ . We say that an event occurs *with high probability* (“w.h.p.”) if its probability tends to one as $n \rightarrow \infty$.

Friedgut [22] proved that for any k there exists a *sharp threshold sequence* $r_{k\text{-NAE}} = r_{k\text{-NAE}}(n)$ such that for any fixed $\varepsilon > 0$ w.h.p. Φ has a NAE-solution if $r < r_{k\text{-NAE}} - \varepsilon$, while w.h.p. Φ fails to have one if $r > r_{k\text{-NAE}} + \varepsilon$. It is widely conjectured but as yet unproven that the threshold sequence converges for any $k \geq 3$. The best previous bounds on $r_{k\text{-NAE}}$ were derived by Achlioptas and Moore [3] and Coja-Oghlan and Zdeborová [12] via the first/second moment method:

$$r_{\text{second}} = 2^{k-1} \ln 2 - \ln 2 + o_k(1) \leq r_{k\text{-NAE}} \leq r_{\text{first}} = 2^{k-1} \ln 2 - \frac{\ln 2}{2} + o_k(1), \quad (1.1)$$

where $o_k(1)$ hides a term that tends to 0 for large k . This left an additive gap of $\frac{1}{2} \ln 2 \approx 0.347$, which our main result closes.

Theorem 1.1. *There is a sequence $\varepsilon_k = 2^{-(1-o_k(1))k}$ such that*

$$2^{k-1} \ln 2 - \left(\frac{\ln 2}{2} + \frac{1}{4}\right) - \varepsilon_k \leq r_{k\text{-NAE}} \leq 2^{k-1} \ln 2 - \left(\frac{\ln 2}{2} + \frac{1}{4}\right) + \varepsilon_k. \quad (1.2)$$

While the numerical improvement obtained in Theorem 1.1 may seem modest, we are going to argue that the result is conceptually quite significant for two reasons. First, we obtain (virtually) matching upper and lower bounds for the first time in a random CSP of this type. Second, and perhaps even more importantly, we devise a rigorous method for taming the condensation phenomenon. Indeed, condensation has been the main obstacle to determining the precise thresholds in random CSPs for the past decade. To understand why, we need to discuss the statistical mechanics picture and its relation to the second moment method.

2 Condensation and the second moment method

The statistical mechanics perspective. We follow [27] to sketch the non-rigorous statistical mechanics approach on random k -NAESAT. Let $\mathcal{S}(\Phi) \subset \{0,1\}^n$ denote the set of NAE-solutions of Φ , and let $Z(\Phi) = |\mathcal{S}(\Phi)|$ be the number of solutions. We turn $\mathcal{S}(\Phi)$ into a graph by considering two solutions σ, τ adjacent if their Hamming distance is $o(n)$. According to [27], the ‘shape’ of $\mathcal{S}(\Phi)$ undergoes two substantial changes w.h.p. at certain densities $0 < r_{\text{sh}} < r_{\text{cond}} < r_{k\text{-NAE}}$.

The first transition occurs at $r_{\text{sh}} \sim 2^{k-1} \ln(k)/k$, almost a factor of k below $r_{k\text{-NAE}}$. Namely, for $r < r_{\text{sh}}$, $\mathcal{S}(\Phi)$ is (essentially) a connected graph. But in the *shattering phase* $r_{\text{sh}} < r < r_{\text{cond}}$, $\mathcal{S}(\Phi)$ splits into connected components $S_1, \dots, S_{N(\Phi)}$ called *clusters* that are mutually separated by a linear Hamming distance $\Omega(n)$. Each cluster S_i only comprises an exponentially small fraction of $\mathcal{S}(\Phi)$. In particular, the total number $N(\Phi)$ of clusters, the so-called *complexity*, is exponential in n . This ‘shattering’ of $\mathcal{S}(\Phi)$ was indeed established rigorously in [1].

As the density r increases beyond r_{sh} , both the overall number $Z(\Phi)$ of solutions and the number and sizes of the clusters shrink. However, the cluster sizes decrease at a slower rate than $Z(\Phi)$, until at density $r_{\text{cond}} = 2^{k-1} \ln 2 - \ln 2 + o_k(1)$ the largest cluster has size $\Omega(Z(\Phi))$ w.h.p. In effect, in the *condensation phase* $r_{\text{cond}} < r < r_{k\text{-NAE}}$, the set $\mathcal{S}(\Phi)$ still decomposes into an exponential number of clusters $S_1, \dots, S_{N(\Phi)}$, each of tiny diameter and all mutually separated by Hamming distance $\Omega(n)$. But in contrast to the shattered phase, now the largest cluster contains a *constant* fraction of the entire set $\mathcal{S}(\Phi)$. Indeed, w.h.p. a *bounded* number of clusters contain a $1 - o(1)$ -fraction of all solutions.

The dominance of a few large clusters in the condensation phase complicates the probabilistic nature of the problem dramatically. To see why, consider the experiment of first choosing a random formula Φ , and then picking two solutions $\sigma, \tau \in \mathcal{S}(\Phi)$ uniformly and independently. For $r_{\text{sh}} < r < r_{\text{cond}}$, σ, τ likely belong to different clusters, and hence can be expected to have a ‘large’ Hamming distance. In fact, it is implicit in the previous work on the second moment method that $\text{dist}(\sigma, \tau) \sim n/2$ w.h.p. [3,12]. Intuitively, this means that the two random solutions ‘decorrelate’. By contrast, for $r_{\text{cond}} < r < r_{k\text{-NAE}}$ both σ, τ belong to the same large cluster with a non-vanishing probability. In effect, with a non-vanishing probability their distance $\text{dist}(\sigma, \tau)$ is tiny, reflecting that solutions in the same cluster are heavily correlated.

The purpose of the physicists’ *Survey Propagation* technique is precisely to deal with this type of correlation. The basic idea is to work with a different, non-uniform probability distribution on $\mathcal{S}(\Phi)$. This *SP distribution* is induced by first choosing a *cluster* S_i uniformly at random among $S_1, \dots, S_{N(\Phi)}$, and then selecting a solution in that cluster S_i uniformly. Since the number $N(\Phi)$ of clusters is (thought to be) exponential in n throughout the condensation phase, two solutions σ', τ' chosen independently from the SP distribution are expected to lie in distinct clusters and thus to decorrelate w.h.p.

Starting from this (appropriately formalized) decorrelation assumption, the SP formalism prescribes a sequence of delicate (non-rigorous) steps to reduce the computation of the *precise* threshold $r_{k\text{-NAE}}$ to the solution of a continuous variational problem for any $k \geq 3$ [14,31]. This variational problem is itself highly

non-trivial, but heuristic numerical techniques yield plausible approximations for small values of k [28]. Moreover, asymptotically for large k the variational problem can be solved analytically. This led to the conjecture that $r_{k\text{-NAE}} = 2^{k-1} \ln 2 - \left(\frac{\ln 2}{2} + \frac{1}{4}\right) + o_k(1)$ [14], which Theorem 1.1 resolves.

Is Theorem 1.1 “optimal”? Of course, it would be interesting to prove that for any k , the *precise* threshold $r_{k\text{-NAE}}$ equals the solution to the variational problem that the SP formalism spits out. However, given that this continuous problem itself appears difficult to solve analytically (to say the very least), it seems that such a result would merely establish the equivalence of two hard mathematical problems. Thus, we believe that Theorem 1.1 marks the end of the line as far as an analytic/explicit computation of $r_{k\text{-NAE}}$ is concerned.

The first and the second moment method. The above statistical mechanics picture holds the key to understanding why the previous arguments did not suffice to pin down $r_{k\text{-NAE}}$ precisely. The best previous bounds (1.1) were obtained by applying the first/second moment method to the number $Z(\Phi)$ of solutions, or a closely related random variable.

With respect to the upper bound, if for some density r the first moment $E[Z(\Phi)]$ tends to 0 as n gets large, then $Z(\Phi) = 0$ w.h.p. by Markov’s inequality. Thus, $r_{k\text{-NAE}} \leq r$. Indeed, it is not difficult to verify that $E[Z(\Phi)] = o(1)$ for $r = r_{\text{first}}$ [3]. This gives the upper bound in (1.1).

The purpose of the second moment method is to bound $r_{k\text{-NAE}}$ from below. The general approach is this: suppose we can define a random variable $Y = Y(\Phi) \geq 0$ such that $Y > 0$ only if Φ has a NAE-solution. Moreover, assume that for some density r , the second moment $E[Y^2]$ satisfies

$$E[Y^2] \leq C \cdot E[Y]^2 \quad (2.1)$$

with $C = C(k) \geq 1$ dependent on k but not on n . Then the *Paley-Zygmund inequality* $P[Y > 0] \geq E[Y]^2 / E[Y^2]$ implies that

$$P[\Phi \text{ has a NAE-solution}] \geq P[Y > 0] \geq E[Y^2] / E[Y]^2 \geq 1/C > 0. \quad (2.2)$$

Because the k -NAESAT threshold is sharp, and as C is independent of n , (2.2) implies that $r_{k\text{-NAE}} \geq r$.

The obvious choice of random variable is the number $Z(\Phi)$ of solutions. Since $Z(\Phi)^2$ is just the number of *pairs* of NAE-solutions, the second moment can be written as

$$E[Z(\Phi)^2] = \sum_{\sigma, \tau \in \{0,1\}^n} P[\text{both } \sigma, \tau \text{ are NAE-solutions}]. \quad (2.3)$$

Indeed, Achlioptas and Moore [3] proved that (2.1) is satisfied for $Y = Z(\Phi)$ if $r \leq 2^{k-1} \ln 2 - (1 + \ln 2)/2$. Improving upon [3], Coja-Oghlan and Zdeborová [12] obtained the best previous lower bound (1.1) by considering a slightly modified random variable $Z'(\Phi)$. Namely, $Z'(\Phi) = Z(\Phi) \cdot \mathbf{1}_{\Phi \in \mathcal{A}}$, where \mathcal{A} is a certain event such that $\Phi \in \mathcal{A}$ w.h.p. In other words, $Z'(\Phi)$ is equal to $Z(\Phi)$ for almost all formulas, but a small fraction of “bad” formulas (that would blow up the second moment) are excluded. Still, $Z'(\Phi)$ admits a similar decomposition as (2.3) (one just has to condition on \mathcal{A}).

As (2.3) shows, the second moment analysis of either $Z(\Phi)$ or $Z'(\Phi)$ boils down to studying the correlations amongst *pairs* of solutions. In fact, it was observed in [3,12] that a *necessary* condition for the success of this approach is that two independently and uniformly chosen $\sigma, \tau \in \mathcal{S}(\Phi)$ satisfy $\text{dist}(\sigma, \tau) \sim n/2$ w.h.p. But according to the statistical mechanics picture, this decorrelation condition is violated for $r > r_{\text{cond}}$ due to the presence of large clusters. Therefore, it is not surprising that the best previous lower bound (1.1) on $r_{k\text{-NAE}}$ coincides with the (conjectured) condensation threshold r_{cond} . Indeed, it was verified in [12] that a certain “weak” form of condensation sets in at $r \sim r_{\text{cond}}$.

The statistical mechanics prescription to overcome these correlations is to work with the Survey Propagation distribution (first select a cluster uniformly, then choose a random solution from that cluster) rather than the uniform distribution over $\mathcal{S}(\Phi)$. This is precisely the key idea behind our new *SP-inspired second moment argument*. Roughly speaking, we are going to develop a way to apply the second moment method

to the number $N(\Phi)$ of *clusters*, rather than the number of solutions. More precisely, we introduce a parameter β that allows us to work with clusters of a prescribed size. A specific choice of β (namely, $\beta = 1/2$) corresponds to the SP distribution and thus to working with $Y(\Phi) = N(\Phi)$.

This new technique allows us to obtain various further results. For instance, we can pin down the typical values of both $Z(\Phi)$ and $N(\Phi)$ throughout the condensation phase (details omitted). Furthermore, our proof entails the following result that confirms the physics conjecture that pairs of solutions drawn from the SP distribution decorrelate throughout the condensation phase.

Corollary 2.1. *Suppose that $r_{\text{cond}} \leq r \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) - \varepsilon_k$. Let σ', τ' be drawn independently from the SP distribution. Then $\text{dist}(\sigma', \tau') = (\frac{1}{2} + o_k(1))n$ w.h.p.*

3 Related work

Rigorous work. The k -NAESAT problem is well-known to be NP-complete in the worst case for any $k \geq 3$. In fact, the NP-complete problem of 2-coloring a k -uniform hypergraph (with $k \geq 3$) simply is the special case of k -NAESAT without negations. The results in [12] are actually phrased in terms of hypergraph 2-coloring but carry over to k -NAESAT directly.

The main contribution of Theorem 1.1 is the improved *lower* bound. In fact, the upper bound in (1.2) can be obtained in several different ways. Achlioptas and Moore [3] state without proof that the (quite intricate) enhanced first moment argument from [16,26] can be used to show that $r_{k\text{-NAE}} \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$. This is indeed plausible as, in terms of the statistical mechanics intuition (which was unknown to the authors of [16,26]) this argument amounts to computing the first moment of the number of *clusters*. Alternatively, generalizing work of Franz and Leone [21], Panchenko and Talagrand [31] proved that the variational problem that results from the SP formalism yields a rigorous upper bound on $r_{k\text{-NAE}}$, which is conjectured to be tight for any $k \geq 3$. The variational problem can be solved asymptotically in the large- k limit (unpublished), yielding the upper bound stated in Theorem 1.1. In this paper we obtain the upper bound by a relatively simple third argument that has a neat combinatorial interpretation.

The proofs of the lower bounds in [3,12] and in the present paper are non-constructive in the sense that they do not entail an efficient algorithm for finding a NAE-solution w.h.p. The best current algorithm for random k -NAESAT is known to succeed for $r \leq O_k(2^k/k)$, a factor of $\Omega_k(k)$ below $r_{k\text{-NAE}}$ [2].

From a statistical mechanics point of view, many random CSPs are similar to random k -NAESAT. In particular, the physics methods suggest the existence of a condensation phase in most random CSPs (e.g., random k -SAT/graph k -coloring). While [3] provided the prototype for the second moment arguments in these and other problems, the technical details in random graph k -coloring [4] or random k -SAT [6] are quite a bit more intricate than in random k -NAESAT.

For instance, random k -NAESAT is simpler than random k -SAT because for any NAE-solution σ the inverse $\bar{\sigma} : x \mapsto 1 - \sigma(x)$ is a NAE-solution as well. This symmetry of the solution space under inversion simplifies the second moment calculations significantly. To cope with the absence of symmetry in random k -SAT, Achlioptas and Peres [6] weighted satisfying assignments cleverly in order to recover the beneficial analytic properties that symmetry induces. Our new second moment method is quite different from this weighting approach, since the asymmetry that called for the weighting scheme in [6] is absent in k -NAESAT.

None of the (few) random CSPs in which the threshold for the existence of solutions is known precisely has a condensation phase. The most prominent example is random k -XORSAT (random linear equations mod 2) [17,32]. In this case, the algebraic nature of the problem precludes condensation: all clusters are simply translations of the kernel. Similarly, the condensation phase is empty in the uniquely extendible problem from [13]. Also in random k -SAT with $k = k(n) > \log_2 n$ (i.e., the clause length grows as a function of n), where the precise threshold has been determined by Frieze and Wormald [23] via the second moment method, condensation does not occur [11]. Nor does it in random 2-SAT [8,25].

Parts of our proof require a precise analysis of geometry of the solution space $\mathcal{S}(\Phi)$. This analysis harnesses some of the ideas that were developed in previous work [1,7,12,15] (e.g., arguments for proving the existence of clusters or of “rigid variables”). However, we need to go beyond these previous arguments significantly in two respects. First, we need to generalize them to accommodate the parameter β that controls the cluster sizes. Second, we need rather precise quantitative information about the cluster structures.

Survey Propagation guided decimation. The SP formalism has given rise to an efficient message passing algorithm called *Survey Propagation guided decimation* (‘SPD’) [29]. Experimentally, SPD seems spectacularly successful at solving, e.g., random k -SAT for small values of k . Unfortunately, no quantitative analysis of this algorithm is currently known (not even a non-rigorous one). The basic idea behind SPD is to approximate the marginals of the SP distribution (i.e., the probability that a given variable is ‘true’ in a solution drawn from the SP distribution) via a message passing heuristic. Then a variable x is selected according to some rule and is assigned a value based on the (approximate) marginal. The entire procedure is repeated on the “decimated” problem instance where x has been eliminated, until (hopefully) a solution is found.

The decorrelation of random solutions chosen from the SP distribution is a crucial assumption behind the message passing computation of the SP marginals. Corollary 2.1 establishes such a decorrelation property rigorously. However, in order to actually analyze SPD, one would have to generalize Corollary 2.1 to the situation of a “decimated” random formula in which a number of variables have already been eliminated by previous steps of the algorithm. Still, we believe that the techniques developed in this paper are a (necessary) first step towards a rigorous analysis of SPD.

4 Heavy solutions and the first moment

In the rest of the paper we sketch the SP-inspired second moment method on which the proof of Theorem 1.1 is based. Aiming for an asymptotic result, we may assume that $k \geq k_0$ for some (large) constant $k_0 > 3$. We also assume $r = 2^{k-1} \ln 2 - \rho$ for some $\frac{1}{2} \ln 2 \leq \rho \leq \ln 2$. Let Φ_i denote the i th clause of the random formula Φ so that $\Phi = \Phi_1 \wedge \dots \wedge \Phi_m$. Furthermore, let Φ_{ij} signify the j th literal of clause Φ_i ; thus, $\Phi_i = \Phi_{i1} \vee \dots \vee \Phi_{ik}$. For a literal ℓ we let $|\ell|$ denote the underlying variable.

As we discussed earlier, the demise of the “standard” second moment method in the condensation phase is due to the dominance of few large clusters. The statistical mechanics prescription for circumventing this issue is to work with a non-uniform distribution over solutions that favors “small” clusters. To implement this strategy, we are going to exhibit a simple parameter that governs the size of the cluster that a solution belongs to. Formally, we define the *cluster* of $\sigma \in \mathcal{S}(\Phi)$ as

$$\mathcal{C}(\sigma) = \mathcal{C}_\Phi(\sigma) = \{\tau \in \mathcal{S}(\Phi) : \text{dist}(\sigma, \tau) \leq 0.01n\}.$$

This definition is vindicated by the following observation from [12], which shows that any two solutions either have the same cluster or are well-separated.

Proposition 4.1. *Suppose that $2^{k-1} \ln 2 - \ln 2 \leq r \leq r_{k-\text{NAE}}$. W.h.p. any two $\sigma, \tau \in \mathcal{S}(\Phi)$ either satisfy $\text{dist}(\sigma, \tau) \leq 0.01n$ or $\text{dist}(\sigma, \tau) \geq (\frac{1}{2} - 2^{-k/3})n$.*

To proceed, we need to get an idea of the “shape” of the clusters $\mathcal{C}(\sigma)$. According to the SP formalism, each cluster has a set $\mathcal{R}(\sigma)$ of $\Omega(n)$ *rigid variables* on which *all* assignments in $\mathcal{C}(\sigma)$ coincide, while the values of the non-rigid variables vary. Formally, we have $\tau(x) = \sigma(x)$ for all $x \in \mathcal{R}(\sigma)$ and all $\tau \in \mathcal{C}(\sigma)$, while for each $x \notin \mathcal{R}(\sigma)$ there is $\tau \in \mathcal{C}(\sigma)$ such that $\tau(x) \neq \sigma(x)$. This implies an immediate bound on the size of $\mathcal{C}(\sigma)$, namely $|\mathcal{C}(\sigma)| \leq 2^{n-|\mathcal{R}(\sigma)|}$. Indeed, we are going to prove that every cluster has a rigid set of size $\Omega(n)$ w.h.p., and that for all clusters w.h.p.

$$\log_2 |\mathcal{C}(\sigma)| = (1 - o_k(1))(n - |\mathcal{R}(\sigma)|). \quad (4.1)$$

With $|\mathcal{C}(\sigma)|$ controlled by the number of rigid variables, it might seem promising to perform first/second moment arguments for the number of solutions with a suitably chosen number of rigid variables. The problem with this is that there is no simple way to tell whether a given variable is rigid: deciding this is NP-hard in the worst case. Intuitively, this is because rigidity emerges from the “global” interplay of variables and clauses. In effect, parametrizing by the number of rigid variables appears technically infeasible.

Instead, we are going to work with a simple “local” parameter that turns out to be a good substitute. Suppose that $x \in \mathcal{R}(\sigma)$. Then x must occur in some clause Φ_i that would be violated if x was assigned the opposite value $1 - \sigma(x)$ (with all other variables unchanged). By the definition of k -NAESAT, this means that the other $k - 1$ literals of Φ_i take the opposite value of the literal whose underlying variable x is. In this case we say that x *supports* Φ_i under σ , and we call Φ_i a *critical* clause. Moreover, we call a variable that supports a clause *blocked*, while all other variables are *free*. While every rigid variable is blocked, the converse is not generally true. Nonetheless, we will see that the number of variables that are blocked but not rigid is small enough so that we can control the cluster sizes in terms of blocked variables.

As a first step, we are going to estimate the expected number of solutions with a given number of blocked variables. Let $\lambda = \frac{kr}{2^{k-1}-1} = k \ln 2 + O_k(k/2^k)$ and let us say that $\sigma \in \mathcal{S}(\Phi)$ is β -heavy if exactly $(1 - \beta) \exp(-\lambda)n$ variables are free. Let $\mathcal{S}_\beta(\Phi)$ be the set of all β -heavy solutions and let $Z_\beta = |\mathcal{S}_\beta(\Phi)|$ denote their number.

Proposition 4.2. *For any $\beta \leq 1$ we have*

$$\mathbb{E}[Z_\beta] = \exp \left[\frac{n}{2^k} (2\rho - \ln(2) - (1 - \beta) \ln(1 - \beta) - \beta + O_k(k \cdot 2^{-k})) \right]. \quad (4.2)$$

In particular, $Z_\beta = 0$ for all $\beta < -3/2$ w.h.p.

Proof. The computation of $\mathbb{E}[Z_\beta]$ is instructive because it hinges upon the solution of an occupancy problem that will play an important role in the second moment computation. Let $\mathbf{1}$ denote the assignment that sets all variables to true. By the linearity of expectation and by symmetry, we have

$$\begin{aligned} \mathbb{E}[Z_\beta] &= \sum_{\sigma \in \{0,1\}^n} \mathbb{P}[\sigma \text{ is a } \beta\text{-heavy solution}] = 2^n \cdot \mathbb{P}[\mathbf{1} \text{ is a } \beta\text{-heavy solution}] \\ &= 2^n \cdot \mathbb{P}[\mathbf{1} \text{ is } \beta\text{-heavy} | \mathbf{1} \text{ is a solution}] \cdot \mathbb{P}[\mathbf{1} \text{ is a solution}]. \end{aligned}$$

Clearly, $\mathbf{1}$ is a solution iff each clause of Φ contains both a positive and a negative literal. A random clause has this property with probability $1 - 2^{1-k}$. Since the $m \sim rn$ clauses are chosen independently, we get

$$2^n \cdot \mathbb{P}[\mathbf{1} \text{ is a solution}] = 2^n (1 - 2^{1-k})^m = \exp \left[\frac{n}{2^k} (2\rho - \ln 2 + O_k(2^{-k})) \right].$$

Working out the conditional probability that $\mathbf{1}$ is β -heavy is not so straightforward. Whether $\mathbf{1}$ is β -heavy depends only on the critical clauses of Φ . Let X be their number. Given that $\mathbf{1}$ is a solution, each clause Φ_i is critical with probability $k/(2^{k-1} - 1)$ independently (as there are $2k$ ways to choose the literal signs to obtain a critical clause). Hence, X has a binomial distribution $\text{Bin}(m, k/(2^{k-1} - 1))$ with mean

$$\mathbb{E}[X | \mathbf{1} \in \mathcal{S}(H)] = \frac{km}{2^{k-1} - 1} = \lambda n.$$

Since the supporting variable of each critical clause is uniformly distributed, given $\mathbf{1} \in \mathcal{S}(H)$ the *expected* number of clauses that each variable supports equals λ . Thinking of the variables as bins and of the critical clauses as balls, standard results on the occupancy problem show that the number of free variables is $(1 + o(1)) \exp(-\lambda)n$ w.h.p. Thus, $\mathbb{E}[Z_\beta]$ is maximized for $\beta = 0$.

By contrast, values $\beta \neq 0$ correspond to *atypical* outcomes of the occupancy problem. Values $\beta < 0$ require an excess number of “empty bins”, while $\beta > 0$ means that fewer bins than expected are empty. To

determine the precise (exponentially small) probability of getting $(1 - \beta) \exp(-\lambda)n$ empty bins, we need to balance large deviations of X against the probability that exactly $(1 - \beta) \exp(-\lambda)n$ bins remain empty for a given value of X . The result of this combined large deviations analysis is the expression (4.2). The analysis also shows that $\mathbb{E}[Z_\beta] = \exp(-\Omega(n))$ for $\beta < -3/2$, whence $Z_\beta = 0$ w.h.p. for $\beta < -3/2$. \square

As a next step, we need to estimate the cluster size of a β -heavy solution.

Proposition 4.3. *W.h.p. for all $-3/2 \leq \beta \leq 1$ all β -heavy $\sigma \in \mathcal{S}(\Phi)$ satisfy*

$$\log_2 |\mathcal{C}(\sigma)| = \frac{n}{2^k} [1 - \beta + o_k(1)]. \quad (4.3)$$

Proof. The crucial thing to show is that all but a very few blocked variables are rigid. The proof of this builds upon arguments developed in [1] to establish rigidity. Suppose that x is blocked in $\sigma \in \mathcal{S}_\beta(\Phi)$, i.e., x supports some clause, say Φ_1 . In any solution τ with $\tau(x) \neq \sigma(x)$ there must be another variable x' that occurs in Φ_1 such that $\tau(x') \neq \sigma(x')$. Given that x supports Φ_1 , the other $k - 1$ variables of Φ_1 are uniformly distributed. Since σ has no more than $(1 - \beta) \exp(-\lambda)n = (1 - \beta + o_k(1))2^{-k}n$ free variables, the probability that x' is free is bounded by $(1 - \beta + o_k(1))(k - 1)/2^k$. In fact, since the *expected* number of clauses that each variable supports is $\lambda = (1 + o_k(1))k \ln 2$, it is quite likely that x' supports several clauses and that therefore “flipping” x' necessitates *several* further flips. Continuing this argument, we see that the number of flips follows a branching process with (initial) successor rate λ . A detailed analysis shows that for all but $O_k(k4^{-k})n$ blocked initial variables x this process will lead to an avalanche of more than $0.01n$ flips, whence $\tau \notin \mathcal{C}(\sigma)$. This shows that all but $o_k(2^{-k})n$ blocked variables are rigid. \square

We are ready to prove that $r_{k-\text{NAE}} \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$, which is (almost) the upper bound promised in Theorem 1.1. (Some additional technical work is needed to replace the $o_k(1)$ by an error term that decays exponentially.) Let $N_\beta = |\{\mathcal{C}(\sigma) : \sigma \in \mathcal{S}(\Phi) \text{ is } \beta\text{-heavy}\}|$ be the number of *clusters* centered around β -heavy solutions. By Proposition 4.3, each such cluster has size $|\mathcal{C}(\sigma)| = 2^{n(1-\beta+o_k(1))/2^k}$ w.h.p. Hence, once more by Proposition 4.3, any solution $\tau \in \mathcal{C}(\sigma)$ is β' -heavy for some β' satisfying $|\beta' - \beta| \leq \delta_k = o_k(1)$ w.h.p. Letting Z_β^* be the total number of β' -heavy solutions with $|\beta' - \beta| \leq \delta_k$, we conclude that

$$N_\beta \cdot 2^{n(1-\beta+o_k(1))/2^k} \leq Z_\beta^* \quad \text{w.h.p.} \quad (4.4)$$

Clearly, $Z_\beta^* \leq \mathbb{E}[Z_\beta^*] \cdot \exp(o(n))$ w.h.p. by Markov’s inequality. Furthermore, as the total number of free variables in each cluster is an integer between 0 and n , we have $\mathbb{E}[Z_\beta^*] \leq (n + 1) \cdot \max_{\beta'} \mathbb{E}[Z_{\beta'}]$. Combining these inequalities with the estimate of $\mathbb{E}[Z_{\beta'}]$ from Proposition 4.2, we find

$$Z_\beta^* \leq \exp[o(n)] \mathbb{E}[Z_\beta^*] \leq \exp\left(\frac{n}{2^k} [2\rho - \ln(2) - (1 - \beta) \ln(1 - \beta) - \beta + o_k(1)]\right) \quad \text{w.h.p.} \quad (4.5)$$

Combining (4.4) and (4.5), we obtain

Fact 4.4. *W.h.p. we have $N_\beta \leq \exp[\eta(\beta) \cdot n/2^k]$ for all β , with*

$$\eta(\beta) = 2\rho - \ln(2) - (1 - \beta) \ln(2 - 2\beta) - \beta + o_k(1). \quad (4.6)$$

Finally, it is a mere exercise in calculus to verify that at density $r^* = 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$ the exponent $\eta(\beta)$ is negative for *all* β . Therefore, Fact 4.4 implies that r^* is an upper bound on $r_{k-\text{NAE}}$.

Remark 5. The exponent $\eta(\beta)$ attains its maximum at $\beta = \frac{1}{2} + o_k(1)$. Together with our second moment bound below, this implies that for $\beta = \frac{1}{2} + o_k(1)$ we have $N(\Phi) = \exp(o_k(1)n) \cdot N_\beta(\Phi)$ w.h.p., i.e., setting $\beta = \frac{1}{2} + o_k(1)$ corresponds to the uniform distribution over clusters and thus to the SP distribution.

5 The second moment

A first attempt. The obvious approach to proving a matching lower bound on $r_{k\text{-NAE}}$ seems to be a second moment argument for the number Z_β of β -heavy solutions, for some suitable β . There is a subtle issue with this, but exploring it will put us on the right track.

We already computed $\mathbb{E}[Z_\beta]$ in Proposition 4.2. As $\mathbb{E}[Z_\beta^2]$ is the expected number of *pairs* of β -heavy solutions, the symmetry properties of the random formula Φ imply that

$$\mathbb{E}[Z_\beta^2] = \mathbb{E}[Z_\beta] \cdot \mathbb{E}[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \quad \text{for any fixed } \sigma \in \{0, 1\}^n.$$

Thus, the second moment condition (2.1) that we would like to establish for $Y = Z_\beta$ becomes

$$\mathbb{E}[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \leq C \cdot \mathbb{E}[Z_\beta]. \quad (5.1)$$

What value of β should we go for? By Fact 4.4 a necessary condition for the existence of β -heavy solutions is that the exponent $\eta(\beta)$ from (4.6) is positive. Let us call β *feasible* for a density r if it is. An elementary calculation shows that for $r > r_{\text{cond}} = 2^{k-1} \ln 2 - \ln 2 + o_k(1)$, any feasible β is strictly positive.

However, (5.1) turns out to be false for *any* $\beta > 0$, for any density $r > 0$. To understand why, let us define the *degree* d_x of a variable $x \in V$ as the number of times that x occurs in the formula Φ . Let $\mathbf{d} = (d_x)_{x \in V}$ be the degree sequence of Φ . It is well known that in the “plain” random formula Φ (without conditioning on $\sigma \in \mathcal{S}_\beta(\Phi)$), the degree of each variable is asymptotically Poisson with mean km/n . On the other hand, if we condition on $\sigma \in \mathcal{S}_\beta(\Phi)$ for some $\beta > 0$, then the degrees are *not* asymptotically Poisson anymore. Indeed, the degree d_x is the sum of the number s_x of clauses that x supports, and the number d'_x of times that x appears otherwise. While d'_x is asymptotically Poisson with mean $< km/n$ as the non-critical clauses do not affect the number of blocked variables at all, s_x is not. More precisely, we saw in the proof of Proposition 4.2 that for $\beta > 0$, s_x is the number of “balls” that x receives in an *atypical* outcome of the occupancy problem. The precise distribution of s_x is quite non-trivial, but it is not difficult to verify that s_x does *not* have a Poisson distribution. Fleshing this observation out leads to the sobering

Lemma 5.1. *For any $\beta > 0$ and any $r > 0$ we have $\mathbb{E}[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \geq \exp(\Omega(n)) \cdot \mathbb{E}[Z_\beta]$.*

In summary, conditioning on $\sigma \in \mathcal{S}_\beta(\Phi)$ with $\beta > 0$ imposes a skewed degree distribution that in turn boosts the expected number of β -heavy solutions beyond the unconditional expectation.

Making things work. We tackle the issue of degree fluctuations by separating the choice of the degree sequence from the choice of the actual formula. More precisely, for a sequence $\mathbf{d} = (d_x)_{x \in V}$ of non-negative integers such that $\sum_{x \in V} d_x = km$ we let $\Phi_{\mathbf{d}}$ denote a k -CNF with degree sequence \mathbf{d} chosen uniformly at random amongst all such formulas. Fixing a “typical” degree sequence \mathbf{d} , we are going to perform a second moment argument for $\Phi_{\mathbf{d}}$, thereby preventing fluctuations of the degrees.

How do we define “typical”? Ideally, we would like \mathbf{d} to enjoy all the properties that the degree sequence of the (unconditioned) random formula Φ is likely to have. Formally, we let $\mathbf{D} = \mathbf{D}_k(n, m)$ be the distribution of the degree sequence of Φ . What we are going to show is that our second moment argument succeeds for a random degree sequence chosen from the distribution \mathbf{D} w.h.p.

Definition 2. A β -heavy solution $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ is good if the following conditions are satisfied.

- We have $|\mathcal{C}(\sigma)| \leq \mathbb{E}[Z_\beta(\Phi_{\mathbf{d}})]$.
- There does not exist $\tau \in \mathcal{S}(\Phi_{\mathbf{d}})$ with $0.01n \leq \text{dist}(\sigma, \tau) \leq (\frac{1}{2} - 2^{-k/3})n$.
- No variable supports more than $3k$ clauses under σ .

The first two items mirror our analysis of the solution space from Section 4. The third one turns out to be useful for a purely technical reason.

Let $\mathcal{S}_{g,\beta}(\Phi_d)$ be the set of good β -heavy solutions and set $Z_{g,\beta}(\Phi_d) = |\mathcal{S}_{g,\beta}(\Phi_d)|$. We perform a second moment argument for $Z_{g,\beta}(\Phi_d)$, with d chosen randomly from the distribution D . The result is

Proposition 5.3. *Suppose that $\beta > 0$ is feasible. There is $C = C(k)$ such that for a degree sequence d chosen from the distribution D w.h.p. $\mathbb{E}[Z_{g,\beta}(\Phi_d)^2] \leq C \cdot \mathbb{E}[Z_{g,\beta}(\Phi_d)]^2$.*

Proposition 5.3 shows that the second moment method for $Z_{g,\beta}(\Phi_d)$ succeeds for feasible β . As we observed in Section 4, a feasible $\beta > 0$ exists so long as $r \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) - O_k(k^4/2^k)$. Hence, Proposition 5.3 and the Paley-Zygmund inequality show that Φ_d is NAE-satisfiable for all such r with a non-vanishing probability for d chosen randomly from D . Consequently, the same is true of the unconditioned formula Φ (because we could generate Φ by first choosing d from D and then generating Φ_d). Since the k -NAESAT threshold is sharp [22], we obtain the lower bound in Theorem 1.1.

Proving Proposition 5.3. As a first step, we need to work out $\mathbb{E}[Z_{g,\beta}(\Phi_d)]$. Suppose $\beta > 0$ is feasible. Recall that ρ is such that $r = 2^{k-1} \ln 2 - \rho$.

Lemma 5.4. *W.h.p. the degree sequence d chosen from D is such that*

$$\mathbb{E}[Z_{g,\beta}(\Phi_d)] \sim \mathbb{E}[Z_\beta(\Phi_d)] = \exp\left[\frac{n}{2^k} (2\rho - \ln 2 - (1 - \beta) \ln(1 - \beta) - \beta + O_k(k/2^k))\right].$$

Proof. Choose and fix a degree sequence d . We need to compute the probability that some $\sigma \in \{0, 1\}^V$ is a good β -heavy solution. By symmetry, we may assume that $\sigma = \mathbf{1}$ is the all-true assignment. Then σ is a solution iff every clause contains both a positive and a negative literal. Since the signs of the literals are chosen for all m clauses independently, we see that

$$\mathbb{P}[\sigma \in \mathcal{S}(\Phi_d)] = (1 - 2^{1-k})^m. \quad (5.2)$$

Given that σ is a solution, the number X of critical clauses has distribution $\text{Bin}(m, k/(2^{k-1} - 1))$, because whether a clause is critical depends on its signs only. As in the proof of Proposition 4.2, to determine the probability that σ is β -heavy we need to solve an occupancy problem: X balls representing the critical clauses are tossed randomly into n bins representing the variables. However, this time the bins have *capacities*: the bin representing $x \in V$ can hold no more than $\min\{3k, d_x\}$ balls in total. Thus, we need to compute the probability that under these constraints, exactly $(1 - \beta)2^{-k}n$ bins are empty. This amounts to a rather non-trivial counting problem, but for a random degree sequence d the probability differs from the formula obtained in Proposition 4.2 only by an error term that decays exponentially in k . More precisely,

$$\mathbb{P}[\sigma \in \mathcal{S}_\beta(\Phi_d) | \sigma \in \mathcal{S}(\Phi_d)] = \exp\left(-\frac{n}{2^k} [(1 - \beta) \ln(1 - \beta) - \beta - O_k(k/2^k)]\right). \quad (5.3)$$

Let us provide some intuition why this is. The bin capacities are such that w.h.p. most bins can hold about $kr = k2^{k-1} \ln 2 + O_k(k)$ balls. By comparison, the total number of balls is $X \sim_k mk/(2^{k-1} - 1) \sim_k n k \ln 2$ w.h.p. In effect, the expected number of balls that a typical bin receives is about $k \ln 2$, way smaller than the capacity of that bin. Indeed, since the number of balls that are received by a typical bin is approximately $\text{Bin}(kr, \frac{nk \ln 2}{km}) \approx \text{Bin}(kr, 2^{-k+1})$, the number of balls can be approximated well by a $\text{Po}(\lambda)$ distribution (with $\lambda = kr/(2^{k-1} - 1) \sim_k k \ln 2$). Thus, the probability that a bin remains empty is close to $\exp(-\lambda)$, which was the probability of the same event in the experiment without capacities. The technical details of this argument are quite delicate, as the fluctuations of the capacities need to be controlled *very* carefully.

Finally, similar arguments as in the proof of Proposition 4.3 yield $\mathbb{P}[\sigma \in \mathcal{S}_{g,\beta}(\Phi_d) | \sigma \in \mathcal{S}_\beta(\Phi_d)] = 1 - o(1)$. Thus, the assertion follows from (5.2)–(5.3). \square

We now turn to the second moment. Fix some $\sigma \in \{0, 1\}^V$, say $\sigma = \mathbf{1}$. Let $Z_{g,\beta}(t, \sigma)$ denote the number of good $\tau \in \mathcal{S}(\Phi_d)$ at distance t from σ . Using the linearity of expectation and recalling that the set of NAE-solutions is symmetric with respect to inversion, we obtain

$$\mathbb{E}[Z_{g,\beta}(\Phi_d) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_d)] \leq 2 \sum_{0 \leq t \leq n/2} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_d)]. \quad (5.4)$$

Let $I = \{t \in \mathbb{Z} : (\frac{1}{2} - 2^{-k/3})n \leq t \leq n/2\}$. The first two conditions from Definition 2 ensure that given that σ is good, with certainty we have

$$\sum_{t \leq 0.01n} Z_{g,\beta}(t, \sigma) \leq |\mathcal{C}(\sigma)| \leq \mathbb{E}[Z_\beta(\Phi_d)] \quad \text{and} \quad \sum_{0.01n < t < (\frac{1}{2} - 2^{-k/3})n} Z_{g,\beta}(t, \sigma) = 0.$$

Hence, Lemma 5.4 and (5.4) yield

$$\mathbb{E}[Z_{g,\beta}(\Phi_d) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_d)] \leq (2 + o(1))\mathbb{E}[Z_{g,\beta}(\Phi_d)] + 2 \sum_{t \in I} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_d)]. \quad (5.5)$$

This reduces the proof to the analysis of the “central terms” with $t \in I$. The result of this is

Lemma 5.5. *There is a constant $C' = C'(k) \geq 1$ such that for a random d we have*

$$\sum_{t \in I} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{\beta,g}(\Phi_d)] \leq C' \cdot \mathbb{E}[Z_{g,\beta}(\Phi_d)] \quad \text{w.h.p.} \quad (5.6)$$

Proof (sketch). This is technically the most challenging bit of this work. The argument boils down to estimating the probability that two random $\sigma, \tau \in \{0, 1\}^n$ with $\text{dist}(\sigma, \tau)/n = \alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ simultaneously are good β -heavy solutions. To compute this probability, we need to analyze the interplay of two occupancy problems as in the proof of Lemma 5.4 with respect to the same degree sequence d .

More precisely, let $B = \bigcup_{x \in V} \{x\} \times \{1, \dots, d_x\}$ be a set of km “balls”. Generating Φ_d is equivalent to drawing a random bijection $\pi : [m] \times [k] \rightarrow B$, with $\pi(i, j) = (x, l)$ indicating that x is the underlying variable of the j th literal of clause i , and independently choosing a map $s : [m] \times [k] \rightarrow \{\pm 1\}$ indicating the signs. Further, we represent the occupancy problems for σ, τ by two “colorings” $g_\sigma, g_\tau : B \rightarrow \{\text{red}, \text{blue}\}$, with $g_\sigma(x, l) = \text{red}$ indicating that the l th position in bin x is occupied under σ (and analogously for τ). We compute the probability $p(\alpha, g_\sigma, g_\tau)$ that π, s induce a formula in which

- literal (i, j) supports clause i under σ iff $g_\sigma \circ \pi(i, j) = \text{red}$, and similarly for τ .
- both σ, τ are good β -heavy solutions.

The result is that for any g_σ, g_τ the “success probability” is *minimized* at $\alpha = 1/2$. Quantitatively,

$$\frac{p(\alpha, g_\sigma, g_\tau)}{p(1/2, g_\sigma, g_\tau)} = \exp \left[O_k(k^4/2^k)(\alpha - 1/2)^2 n \right] \quad \text{for any } g_\sigma, g_\tau. \quad (5.7)$$

On the other hand, the total *number* of assignment pairs satisfies

$$\frac{|\{(\sigma, \tau) : \text{dist}(\sigma, \tau) = \alpha n\}|}{|\{(\sigma, \tau) : \text{dist}(\sigma, \tau) = n/2\}|} = \binom{n}{\alpha n} / \binom{n}{n/2} = \exp(-(4 - o_k(1))(\alpha - 1/2)^2 n), \quad (5.8)$$

which is *maximized* at $\alpha = 1/2$. Combining (5.7) and (5.8), we see that for any two colorings g_σ, g_τ the dominant contribution to the second moment stems from $\alpha = \frac{1}{2} + O(1/\sqrt{n})$, i.e., from “perfectly decorrelated” σ, τ . The assertion follows by evaluating the contribution of such α explicitly and summing over g_σ, g_τ . \square

Acknowledgment. The first author thanks Dimitris Achlioptas and Lenka Zdeborová for helpful discussions on the second moment method and the statistical mechanics work on random CSPs.

References

1. D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.
2. D. Achlioptas, J.H. Kim, M. Krivelevich, P. Tetali: Two-coloring random hypergraphs. Random Structures and Algorithms **18** (2002), 249–259.
3. D. Achlioptas, C. Moore: Random k -SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
4. D. Achlioptas, A. Naor: The two possible values of the chromatic number of a random graph. Annals of Mathematics **162** (2005) 1333–1349.
5. D. Achlioptas, A. Naor, Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.
6. D. Achlioptas, Y. Peres: The threshold for random k -SAT is $2^k \ln 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.
7. D. Achlioptas, F. Ricci-Tersenghi: On the solution space geometry of random constraint satisfaction problems. Proc. 38th STOC (2006) 130–139.
8. V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.
9. A. Coja-Oghlan: A better algorithm for random k -SAT. SIAM J. Computing **39** (2010) 2823–2864.
10. A. Coja-Oghlan, C. Efthymiou: On independent sets in random graphs. Proc. 22nd SODA (2011) 136–144.
11. A. Coja-Oghlan, A. Frieze: Random k -SAT: the limiting probability for satisfiability for moderately growing k . Electronic Journal of Combinatorics **15** (2008) N2.
12. A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. Proc. 23rd SODA (2012), to appear.
13. H. Connamacher, M. Molloy: The exact satisfiability threshold for a potentially intractable random constraint satisfaction problem. Proc. 45th FOCS (2004) 590–599.
14. L. Dall’Asta, A. Ramezani, R. Zecchina: Entropy landscape and non-Gibbs solutions in constraint satisfaction problems. Phys. Rev. E **77**, 031118 (2008).
15. H. Daudé, M. Mézard, T. Mora, R. Zecchina: Pairs of SAT-assignments in random Boolean formulae. Theoretical Computer Science **393** (2008) 260–279.
16. O. Dubois, Y. Boufkhad: A general upper bound for the satisfiability threshold of random r -SAT formulae. J. Algorithms **24** (1997) 395–420.
17. O. Dubois, J. Mandler: The 3-XORSAT threshold. Proc. 43rd FOCS (2002) 769–778.
18. C. Efthymiou: A simple algorithm for random colouring $G(n, d/n)$ using $(2+\epsilon)d$ colours. Proc. 23rd SODA (2012), to appear.
19. U. Feige, E. Mossel, D. Vilenchik: Complete convergence of message passing algorithms for some satisfiability problems. Proc. 10th RANDOM (2006) 339–350.
20. P. Flajolet, R. Sedgewick: Analytic Combinatorics. Cambridge University Press, Cambridge, 2009.
21. S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. J. Statist. Phys. **111** (2003) 535–564.
22. E. Friedgut: Hunting for sharp thresholds. Random Struct. Algorithms **26** (2005) 37–51.
23. A. Frieze, N. Wormald: Random k -Sat: a tight threshold for moderately growing k . Combinatorica **25** (2005) 297–305.
24. A. Gerschenfeld, A. Montanari: Reconstruction for models on random graphs. Proc. 48th FOCS (2007) 194–204.
25. A. Goerdts: A threshold for unsatisfiability. Proc. 17th MFCS (1992) 264–274.
26. L. Kirovski, E. Kranakis, D. Krizanc, Y. Stamatou: Approximating the unsatisfiability threshold of random formulas. Random Structures Algorithms **12** (1998) 253–269.
27. F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.
28. S. Mertens, M. Mézard, R. Zecchina: Threshold values of random K -SAT from the cavity method. Random Struct. Alg. **28** (2006) 340–373.
29. M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.
30. A. Montanari, R. Restrepo, P. Tetali: Reconstruction and clustering in random constraint satisfaction problems. SIAM J. Discrete Math. **25** (2011) 771–808.
31. D. Panchenko, M. Talagrand: Bounds for diluted mean-fields spin glass models. Probab. Theory Relat. Fields **130** (2004) 319–336.
32. B. Pittel, G. Sorkin: The satisfiability threshold for k -XORSAT. Preprint (2011).

Appendix

This appendix contains the details omitted from the extended abstract. Section A contains some preliminary facts about random variables that will be used many times. Appendix B contains the full proof of the upper bound claimed in Theorem 1.1 (with ε_k exponentially small in k). Finally, in Appendices C and D we carry out the second moment argument in full.

A Preliminaries

The next lemma provides an asymptotically tight bound for the probability that a sum of independent and identically distributed random variables attains a specific value. It will be an important tool in our further analysis, since we will be often interested in the exact probabilities of *rare* events.

Lemma A.1. *Let X_1, \dots, X_n be independent random variables with support on \mathbf{N}_0 with probability generating function $P(z)$. Let $\mu = \mathbb{E}[X_1]$ and $\sigma^2 = \text{Var}[X_1]$. Assume that $P(z)$ is an entire and aperiodic function. Then, uniformly for all $T_0 < \alpha < T_\infty$, where $T_x = \lim_{z \rightarrow x} \frac{zP'(z)}{P(z)}$, as $n \rightarrow \infty$*

$$\Pr[X_1 + \dots + X_n = \alpha n] = (1 + o(1)) \frac{1}{\zeta \sqrt{2\pi n \xi}} \left(\frac{P(\zeta)}{\zeta^\alpha} \right)^n, \quad (\text{A.1})$$

where ζ and ξ are the solutions to the equations

$$\frac{\zeta P'(\zeta)}{P(\zeta)} = \alpha \quad \text{and} \quad \xi = \frac{d^2}{dz^2} (\ln P(z) - \alpha \ln z) \Big|_{z=\zeta}. \quad (\text{A.2})$$

Moreover, there is a $\delta_0 > 0$ such that for all $0 \leq |\delta| \leq \delta_0$ the following holds. If $\alpha = \mathbb{E}[X_1] + \delta\sigma$, then

$$\Pr[X_1 + \dots + X_n = \alpha n] = (1 + O(\delta)) \frac{1}{\sqrt{2\pi n \sigma}} e^{(-\delta^2/2 + O(\delta^3))n}. \quad (\text{A.3})$$

Proof. The first statement follows immediately from Theorem VIII.8 and the remark after Example VIII.11 in [20]. To see the second statement let us write ζ_δ for the solution to the equation $\frac{\zeta_\delta P'(\zeta_\delta)}{P(\zeta_\delta)} = \mu + \delta\sigma$. Since $P(1) = 1$ and $P'(1) = \mu$ we infer that if $\delta = 0$, then $\zeta_\delta = 1$. Moreover, a Taylor series expansion around $z = 1$ guarantees for all δ in a bounded interval around 0 that

$$\mu + \delta\sigma = \frac{\zeta_\delta P'(\zeta_\delta)}{P(\zeta_\delta)} = \frac{P'(1)}{P(1)} + (\zeta_\delta - 1) \frac{P''(1) + P'(1) - \frac{P'(1)^2}{P(1)}}{P(1)} + O((\zeta_\delta - 1)^2).$$

Since $\sigma^2 = P''(1) + P'(1) - P'(1)^2$, for all δ in a bounded interval around 0 we have that $\zeta_\delta = 1 + \delta/\sigma + O(\delta^2)$. In order to show (A.3) we evaluate the right-hand side of (A.1) at $\zeta = \zeta_\delta$. Again a Taylor series expansion around $z = 1$ guarantees that

$$\begin{aligned} \frac{P(\zeta_\delta)}{\zeta_\delta^\alpha} &= P(1) + (\zeta_\delta - 1)(P'(1) - \alpha P(1)) + \frac{(\zeta_\delta - 1)^2}{2} (P''(1) + P(1)\alpha^2 + P(1)\alpha - 2P'(1)\alpha) + O(\delta^3) \\ &\stackrel{(\alpha=\mu+\delta)}{=} 1 - \delta^2 + \frac{\delta^2}{2\sigma^2} (P''(1) + \mu - \mu^2 + O(\delta)) + O(\delta^3) \\ &= 1 - \frac{\delta^2}{2} + O(\delta^3). \end{aligned}$$

The exponential term in (A.3) is then obtained by using the fact $1 - x = e^{-x - \Theta(x^2)}$. Finally, note that

$$\frac{d^2}{dz^2} (\ln P(z) - \alpha \ln z) = \frac{P''(z)}{P(z)} - \frac{P'(z)^2}{P(z)^2} + \frac{\alpha}{z^2}.$$

By applying again Taylor's Theorem to this function we obtain after some elementary algebra (details omitted) that the value of this function at $\zeta = \zeta_\delta$ equals $\sigma + O(\delta)$, and the proof of (A.3) is completed. \square

The next statement provides tight asymptotic bounds for binomial coefficients.

Proposition A.2. *Let $0 < \alpha \leq 1/2$ and $-1/2 < \varepsilon < 1/2$ be such that $0 < \alpha + \varepsilon < 1$. Then, as $N \rightarrow \infty$*

$$\binom{N}{\alpha N} = \frac{1 + o(1)}{\sqrt{2\pi f(\alpha)N}} e^{H(\alpha)N} \quad \text{and} \quad \binom{N}{(\alpha + \varepsilon)N} = \frac{1 + o(1)}{\sqrt{2\pi f(\alpha + \varepsilon)N}} e^{(H(\alpha) + \varepsilon \log(\frac{1-\alpha}{\alpha}) + O(\varepsilon^2/\alpha))N},$$

where $H(x) = -x \ln x - (1-x) \ln(1-x)$ denotes the entropy function and $f(x) = x(1-x)$.

Proof. The first statement is well-known, see e.g. [20]. To see the second statement, note first that that $H'(x) = \ln(\frac{1-x}{x})$ and $H''(x) = (x(x-1))^{-1}$, both valid in $(0, 1)$. Then, Taylor's Theorem guarantees that

$$H(\alpha + \varepsilon) = H(\alpha) + \varepsilon H'(\alpha) + O(\varepsilon^2/\alpha),$$

from which the second statement follows immediately. \square

B The upper bound on $r_{k-\text{NAE}}$

To prove the upper bound on $r_{k-\text{NAE}}$ we are going to combine the upper bound on the expectation of Z_β from Proposition 4.2 with a *lower* bound on the cluster sizes of β -heavy assignments, see Lemma B.3. Let $\lambda = kr/(2^{k-1} - 1)$. First of all, we fill the missing pieces in the proof of Proposition 4.2. The next lemma provides the analysis for the balls-into-bins game that was omitted in the proof of Proposition 4.2.

Lemma B.1. *Let $X \sim \text{Bin}(m, k/(2^{k-1} - 1))$. We throw X balls into n bins uniformly at random. Let B_i denote the number of bins that receive i balls. Then, for any $-3/2 \leq \beta \leq 1$*

$$n^{-1} \ln \Pr [B_0 = (1 - \beta)e^{-\lambda}n] = n^{-1} \ln \Pr [\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n] + O_k(k4^{-k}). \quad (\text{B.1})$$

Proof. We shall estimate the desired probability by conditioning on any specific value x of X . Let F_i be the number of balls in the i th bin, and let P_1, \dots, P_n be independent Poisson distributed random variables with mean λ . It is well-known and easy to verify that the distribution of (F_1, \dots, F_n) is the same as the distribution of (P_1, \dots, P_n) , *conditioned on the event* $\mathcal{A}(x) = \sum_{1 \leq i \leq n} P_i = x$. So, if we denote by N_0 the number of P_i 's that are equal to 0, we infer that

$$\Pr [B_0 = (1 - \beta)e^{-\lambda}n \mid X = x] = \Pr [N_0 = (1 - \beta)e^{-\lambda}n \mid \mathcal{A}(x)].$$

By the law of total probability this equals

$$\Pr [B_0 = (1 - \beta)e^{-\lambda}n \mid X = x] = \Pr [N_0 = (1 - \beta)e^{-\lambda}n] \cdot \frac{\Pr[\mathcal{A}(x) \mid N_0 = (1 - \beta)e^{-\lambda}n]}{\Pr[\mathcal{A}(x)]}.$$

Note that $N_0 \sim \text{Bin}(n, e^{-\lambda})$. Furthermore, if we denote by $P'_1, \dots, P'_{\xi n}$, where $\xi = 1 - (1 - \beta)e^{-\lambda}$, independent Poisson variables that are conditioned on being at least 1, then the above equation implies that

$$\frac{\Pr [B_0 = (1 - \beta)e^{-\lambda}n]}{\Pr [\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n]} = \sum_{x=\xi n}^m \frac{\Pr[\sum_{i=1}^{\xi n} P'_i = x]}{\Pr[\text{Po}(\lambda n) = x]} \cdot \Pr [\text{Bin}(rn, k/(2^{k-1} - 1)) = x]. \quad (\text{B.2})$$

In order to complete the proof of (B.1) we will derive in the sequel appropriate bounds for the right-hand side of the above equation. First, to obtain a lower bound, note that $\xi < \lambda$, since $\xi < 1$ and $\lambda = k \ln 2 + O_k(k2^{-k})$, which is > 1 for sufficiently large k . Thus, we can obtain a lower bound for (B.2) by considering only the term in the sum that corresponds to $x = \lambda n$. Since $E[\text{Po}(\lambda n)] = E[\text{Bin}(rn, k/(2^{k-1} - 1))] = \lambda n$, we infer by applying Lemma A.1 that

$$\Pr[\text{Po}(\lambda n) = \lambda n] = \Theta(n^{-1/2}) \quad \text{and} \quad \Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \lambda n] = \Theta(n^{-1/2}).$$

It remains to bound $\Pr[\sum_{i=1}^{\xi n} P'_i = \lambda n]$. Note that $E[P'_1] = \frac{\lambda}{1-e^{-\lambda}}$. If we write $N = \xi n$, then

$$\Pr \left[\sum_{i=1}^{\xi n} P'_i = \lambda n \right] = \Pr \left[\sum_{i=1}^N P'_i = \left(E[X_1] + \frac{\beta \lambda e^{-\lambda}}{\xi(1-e^{-\lambda})} \right) N \right],$$

i.e., we require that the sum of the P'_i 's deviates from the expected value by $O_k(k2^{-k}n)$. By applying Lemma A.1, where we set $\delta = O_k(k^{1/2}2^{-k})$, we conclude that the right-hand side of (B.2) is at least $\exp\{-O_k(k4^{-k}n)\}$. This shows the lower bound in (B.1).

In the remainder of this proof we will show an upper bound for the right-hand side of (B.2). To this end, we will argue that the ratio $\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n] / \Pr[\text{Po}(\lambda n) = \gamma \lambda n]$ is essentially bounded for all x in the given range, from which the claim immediately follows. More specifically, let us write $x = \gamma \lambda n$, where $\xi/\lambda \leq \gamma \leq r/\lambda$. By applying Stirling's Formula $N! = (1 + o(1))\sqrt{2\pi N}(N/e)^N$ we infer that

$$\Pr[\text{Po}(\lambda n) = \gamma \lambda n] = \Theta(1) n^{-1/2} \exp\{\lambda n(-1 + \gamma - \gamma \ln \gamma)\}. \quad (\text{B.3})$$

Moreover, by abbreviating $p = k/(2^{k-1} - 1)$ we get

$$\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n] = \binom{rn}{(\gamma p) rn} p^{(\gamma p) rn} (1-p)^{(1-\gamma p) rn}.$$

Since $\binom{N}{\alpha N} \leq e^{H(\alpha)N}$, where H denotes the entropy function, we obtain after some elementary algebra

$$\Pr[\text{Bin}(rn, p) = \gamma \lambda n] \leq \exp \left\{ \lambda n \left(-\gamma \ln \gamma - \frac{1-\gamma p}{p} \ln \left(\frac{1-\gamma p}{1-p} \right) \right) \right\}.$$

By combining this with (B.3) we obtain the estimate

$$\frac{\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n]}{\Pr[\text{Po}(\lambda n) = \gamma \lambda n]} \leq \Theta(\sqrt{n}) e^{f(\gamma) \lambda n}, \quad \text{where} \quad f(\gamma) = 1 - \gamma - \frac{1-\gamma p}{p} \ln \left(\frac{1-\gamma p}{1-p} \right).$$

Recall that $0 < \xi/\lambda \leq \gamma \leq r/\lambda = 1/p$, and note that both $f(0)$ and $f(1/p)$ are < 0 . Moreover, f has an extremal point at $\gamma = 1$, where $f(1) = 0$. Thus, for all γ in the considered range we have that $f(\gamma) \leq 0$, which implies that the right-hand side of (B.2) is bounded from above by at most a polynomial in n . This completes the proof of the lemma. \square

The proof of Proposition 4.2 then completes by applying the following statement.

Lemma B.2. *There is a $k_0 \geq 3$ such that the following is true. Let $Y \sim \text{Bin}(n, e^{-\lambda})$. For any $-3/2 \leq \beta \leq 1$*

$$n^{-1} \ln \Pr \left[Y = \lfloor (1 - \beta)e^{-\lambda} n \rfloor \right] = f(\beta) + O_k(4^{-k}).$$

Proof. Let us abbreviate $\xi = (1 - \beta)e^{-\lambda}$. We will assume that $\xi n = \lfloor \xi n \rfloor$, i.e., that $\beta = 1 - N(e^{-\lambda}n)^{-1}$ for some $N \in \mathbf{N}_0$. To see that this is sufficient, note that by Taylor's Theorem, for any $\beta \geq 1$ and any $|\varepsilon_n| \leq (e^{-\lambda}n)^{-1}$ such that $\beta + \varepsilon_n \leq 1$ there is a $\delta \in [\beta, \beta + \varepsilon_n]$ such that

$$f(\beta + \varepsilon_n) = f(\beta) + \varepsilon_n f'(\delta) = f(\beta) + \varepsilon_n e^{-\lambda} \ln(1 - \delta) = f(\beta) + O_k(4^{-k}).$$

With the above assumption we proceed with the proof of the claim. The definition of the binomial distribution implies

$$\Pr[Y = (1 - \beta)e^{-\lambda}n] = \binom{n}{\xi n} e^{-\lambda \xi n} (1 - e^{-\lambda})^{(1-\xi)n}. \quad (\text{B.4})$$

If $\beta = 1$, then $\xi = 0$ the above expression simplifies to

$$(1 - e^{-\lambda})^n = \exp\{n \ln(1 - e^{-\lambda})\} = \exp\{n(-e^{-\lambda} - \Theta(e^{-2\lambda}))\}.$$

Since $f(1) = e^{-\lambda}$ and $\lambda = k \ln 2 + \Theta(k2^{-k})$, we infer that the statement is true for $\beta = 1$. It remains to treat the case $\beta < 1$. Standard bounds for the binomial coefficients imply

$$\binom{n}{\xi n} = \frac{\Theta(1)}{\sqrt{\xi(1-\xi)n}} e^{nH(\xi)}, \quad \text{where} \quad H(x) = -x \ln x - (1-x) \ln(1-x).$$

Using the estimate $\ln(1-x) = -x - \Theta(x^2)$, which is valid for $|x| < 1$, we infer after some elementary algebra that

$$n^{-1} \ln \binom{n}{\xi n} = e^{-\lambda}((1-\beta)\lambda - (1-\beta) \log(1-\beta) + (1-\beta)) + O_k(4^{-k}) \quad (\text{B.5})$$

Similarly, the second and the third term in (B.4) can be estimated with

$$n^{-1} \ln (e^{-\lambda \xi n} (1 - e^{-\lambda})^{(1-\xi)n}) = -e^{-\lambda}((1-\beta)\lambda + 1) + O_k(4^{-k}).$$

By plugging this fact together with (B.5) into (B.4) we finally obtain the desired statement. \square

We proceed with the proof of the upper bound in Theorem 1.1. Let $Z_{\beta, \gamma}$ denote the number of β -heavy solutions σ such that $\frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \leq (1 - \beta - \gamma)e^{-\lambda}$. The following statement provides an upper bound for the expected number of such solutions.

Lemma B.3. *For any $-3/2 \leq \beta \leq 1$ and $\gamma > k^{5/2}e^{-\lambda}$ we have for sufficiently large k*

$$\frac{1}{n} \ln \mathbb{E}[Z_{\beta, \gamma}] \leq \frac{1}{n} \ln \mathbb{E}[Z_{\beta}] - \frac{\ln k}{6} \gamma e^{-\lambda}.$$

Proof. Let $\sigma \in \{0, 1\}^V$ be an assignment; for the sake of concreteness, assume that $\sigma = \mathbf{1}$. In order to bound $\mathbb{E}[Z_{\beta, \gamma}]$ it is sufficient to estimate the probability of the event

$$\mathcal{E} = \left\{ \frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \leq (1 - \beta - \gamma)e^{-\lambda} \right\},$$

given that $\sigma \in \mathcal{S}_{\beta}(\Phi)$. Let $\mathcal{F}(\sigma)$ denote the set of free variables, and denote by \mathcal{X} be the set of clauses that do not contain both a positive and a negative literal whose underlying variable is in $V \setminus \mathcal{F}(\sigma)$. Then only the clauses in \mathcal{X} impose constraints on the free variables. We decompose \mathcal{X} into $k-1$ subsets $\mathcal{X}_2, \dots, \mathcal{X}_k$, where \mathcal{X}_i the set of all clauses in \mathcal{X} that contain i variables from $\mathcal{F}(\sigma)$. Note that $\mathcal{X} = \cup_{i=2}^k \mathcal{X}_i$, as any clause with only one variable from $\mathcal{F}(\sigma)$ necessarily contains both positive and negative literals whose underlying

variables are not free. Let $X_i = |X_i|$. Since only the clauses in \mathcal{X} impose constraints on variables from $\mathcal{F}(\sigma)$ that occur in them, we infer that

$$\frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \geq |\mathcal{F}(\sigma)| - Y, \text{ where } Y = \sum_{i=2}^k iX_i.$$

In the remainder we will show that

$$\frac{1}{n} \ln \Pr \left[Y > \gamma e^{-\lambda} n \mid \sigma \in \mathcal{S}_\beta(\Phi) \right] \leq -\frac{\ln k}{6} \gamma e^{-\lambda}, \quad (\text{B.6})$$

from which the statement in the lemma follows immediately.

Note that the set $\mathcal{F}(\sigma)$ is determined by the critical clauses only. Therefore, given that $\sigma \in \mathcal{S}_\beta(\Phi)$, the variables that occur in the non-critical clauses are independent and uniformly distributed over the set of all variables. Similarly, given that $\sigma \in \mathcal{S}_\beta(\Phi)$ the $k-1$ variables that contributed the “majority value” to each critical clause are independently uniformly distributed. Therefore, X_i is stochastically dominated by a binomial random variable

$$X'_i \sim \text{Bin}(m, p_i), \text{ where } p_i = 2^{-k+1} \cdot 2^i \binom{k}{i} ((1-\beta) \exp(-\lambda))^i.$$

Our assumption $-3/2 \leq \beta \leq 1$ guarantees that $(1-\beta)e^{-\lambda} \leq 3e^{-\lambda} \leq 3 \cdot 2^{-k}$. By using the estimate $\binom{k}{i} \leq k^i$ we infer that

$$p_i \leq 2^{-k+1} \cdot 2^i \binom{k}{i} ((1-\beta)e^{-\lambda})^i \leq 2^{-k+1} (6k2^{-k})^i. \quad (\text{B.7})$$

Moreover, note that the X_i are negatively correlated. Indeed, let $X_{i,j}$ be the indicator for the event that the clause $\Phi_j \in \mathcal{X}_i$. Then, for all $i \neq i'$ we have $\mathbb{E}[X_{i,j} X_{i',j}] = 0 \leq \mathbb{E}[X_{i,j}] \mathbb{E}[X_{i',j}]$, and otherwise, if $(i, j) \neq (i', j')$, then $X_{i,j}$ and $X_{i',j'}$ are independent. Thus, for any $\delta > 0$, Markov's inequality implies with $t = \gamma e^{-\lambda}$

$$\Pr[Y > t \mid \sigma \in \mathcal{S}_\beta(\Phi)] \leq e^{-\delta t} \prod_{i=2}^k \mathbb{E}[e^{\delta i X_i}] \leq e^{-\delta t} \prod_{i=2}^k \mathbb{E}[e^{\delta i X'_i}] \leq e^{-\delta t} \prod_{i=2}^k (p_i e^{\delta i} + 1 - p_i)^m,$$

Let us fix $\delta = \frac{1}{5} \ln k$. By the arithmetic-geometric mean inequality we obtain that the expression in the previous equation is at most

$$e^{-\delta t} \left(\frac{\sum_{i=2}^k p_i e^{\delta i} + 1 - p_i}{k} \right)^{km} \stackrel{(\text{B.7})}{\leq} e^{-\delta t} \left(1 + \frac{2^{-k+1} \sum_{i=2}^k (6k2^{-k})^i e^{\delta i}}{k} \right)^{km} = e^{-\delta t} e^{O_k(k^2 4^{-k} e^{2\delta})n}.$$

Since $t = \gamma e^{-\lambda} > k^{5/2} 4^{-k}$, for sufficiently large k we get (B.6), and the proof is completed. \square

Consider the function

$$g(\beta) = h(\beta) - (1-\beta)e^{-\lambda} \ln 2,$$

where

$$h(\beta) = \frac{2\rho - \ln 2}{2^k} + f(\beta) \quad \text{and} \quad f(\beta) = -((1-\beta) \ln(1-\beta) + \beta) e^{-\lambda}.$$

Let r_* be the least density r such that $g(\beta) < -k^3 4^{-k+1}$ for all $\beta \geq -1$. Since g is maximized for $\beta = 1/2$, where $g(1/2) = \frac{2\rho - \ln 2}{2^k} - \frac{1}{2} e^{-\lambda}$, it is easily verified that

$$r_* = 2^{k-1} \ln 2 - \left(\frac{\ln 2}{2} + \frac{1}{4} \right) + O_k(k^3 2^{-k}).$$

Proposition B.4. *With $r = r^*$ the random formula Φ does not have a NAE-solution w.h.p.*

Proof. Let $Z_{\leq \beta}$ be the number of solutions that are β' -heavy for some $\beta' \leq \beta$. In order to prove that $Z_{\leq \beta} = 0$ w.h.p. for all β we proceed as follows. Let $-3/2 = \beta_0 < \dots < \beta_\ell = 1$ be a sequence such that $|\beta_i - \beta_{i+1}| \leq \delta$ for all i , where $\delta = 2^{-2^k}$. We are going to show inductively that $Z_{\leq \beta_i} = 0$ w.h.p.; by the previous discussion we may assume that this is true for $i = 0$.

Let us assume for the induction step that i is such that w.h.p. $Z_{\leq \beta_i} = 0$. Let $\gamma_0 = k^3 e^{-\lambda}$, and let Z' be the number of solutions that are β' -heavy for some $\beta' > \beta_i$ and such that $\frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \geq (1 - \beta_i - \gamma_0) e^{-\lambda}$. Then, by applying Proposition 4.2 and using that $h(x)$ is monotone increasing for $x \leq 0$ and monotone decreasing for $x \geq 0$ we obtain

$$\frac{1}{n} \ln \mathbb{E}[Z'] \leq \max_{\beta > \beta_i} h(\beta) + O_k(\delta + k4^{-k}) = O_k(k4^{-k}) + \begin{cases} h(0), & \text{if } \beta_i \leq 0, \\ h(\beta_i), & \text{if } \beta_i > 0 \end{cases}.$$

Let us first consider the case $\beta_i \leq 0$. The choice of r^* guarantees that $g(0) = h(0) - e^{-\lambda} \ln 2 < -k^3 4^{-k+1}$. Since $Z' > 0$ implies $Z' \geq \exp\{n(1 - \beta_i - \gamma_0) e^{-\lambda} \ln 2\} \geq \exp\{n(1 - \gamma_0) e^{-\lambda} \ln 2\}$ or otherwise $Z_{\leq \beta_i} > 0$ we infer for sufficiently large k that

$$\Pr[Z' > 0] \leq \Pr[Z_{\leq \beta_i} > 0] + \mathbb{E}[Z'] \exp\{-n(1 - \gamma_0) e^{-\lambda} \ln 2\} = o(1).$$

On the other hand, if $\beta_i > 0$, then again the choice of r^* is such that $g(\beta_i) = h(\beta_i) - (1 - \beta_i) e^{-\lambda} \ln 2 < -k^3 4^{-k+1}$. Thus, for sufficiently large k

$$\frac{1}{n} \ln \mathbb{E}[Z'] < -k^3 4^{-k+1} + (1 - \beta_i) e^{-\lambda} \ln 2 + O_k(k4^{-k}) < -k^7 4^{-k} + (1 - \beta_i - \gamma_0) e^{-\lambda} \ln 2.$$

So, since $Z' > 0$ implies $Z' \geq \exp\{n(1 - \beta_i - \gamma_0) e^{-\lambda} \ln 2\}$ or otherwise $Z_{\leq \beta_i} > 0$ we infer that

$$\Pr[Z' > 0] \leq \Pr[Z_{\leq \beta_i} > 0] + \mathbb{E}[Z'] \exp\{-n(1 - \beta_i - \gamma_0) e^{-\lambda} \ln 2\} = o(1).$$

Thus, in both cases we have that $\Pr[Z' > 0] = o(1)$. It remains to consider all satisfying assignments such that $\frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \leq (1 - \beta_i - \gamma_0) e^{-\lambda}$. More specifically, let Z'_j be the number of solutions that are β' -heavy for some $\beta_i < \beta' \leq \beta_{i+1}$ and such that

$$(1 - \beta_i - \gamma_{j+1}) e^{-\lambda} \leq \frac{1}{n} \log_2 |\mathcal{C}(\sigma)| \leq (1 - \beta_i - \gamma_j) e^{-\lambda},$$

where $\gamma_{j+1} = 2\gamma_j$. Choose β' be such that $\mathcal{S}_{\beta'}(\Phi) \cap \mathcal{C}(\sigma)$ is maximized. Then

$$|\mathcal{S}_{\beta'}(\Phi) \cap \mathcal{C}(\sigma)| \geq \frac{|\mathcal{C}(\sigma)|}{n}. \tag{B.8}$$

Since $Z_{\leq \beta_i} = 0$ w.h.p., we may assume that $\beta' > \beta_i$. There are two cases to consider.

Case 1: $1 - \beta_i - \gamma_{j+1} > 1 - \beta'$. We will show that in this case the number of β' -heavy assignments is larger than the expected value by at least an exponential factor. Indeed, our assumption on g implies for sufficiently large k that

$$\frac{1}{n} \ln \mathbb{E}[Z_{\beta'}] = h(\beta') + O_k(k4^{-k}) < -k^3 4^{-k} + (1 - \beta') e^{-\lambda} \ln 2 < -k^3 4^{-k} + (1 - \beta_i - \gamma_{j+1}) e^{-\lambda} \ln 2.$$

However, if (B.8) holds then

$$\frac{1}{n} \ln Z_{\beta'} \geq \frac{1}{n} \ln |\mathcal{C}(\sigma)| - o(1) = (1 - \beta_i - \gamma_{j+1}) e^{-\lambda} \ln 2 - o(1).$$

By Markov's inequality, the probability of this event is $\exp(-\Omega(n))$.

Case 2: $1 - \beta_i - \gamma_{j+1} \leq 1 - \beta'$. The assumption guarantees the existence of a $\gamma' > 0$ such that

$$1 - \beta_i - \gamma_{j+1} = 1 - \beta' - \gamma'.$$

In this case we will show that the number of solutions in $\mathcal{S}_{\beta', \gamma'}(\Phi)$ is larger than the expected value by at least an exponential factor. Equation (B.8) implies that

$$\frac{1}{n} \ln \mathbb{E} [Z_{\beta', \gamma'}] \geq \frac{1}{n} \ln |\mathcal{C}(\sigma)| - o(1) = (1 - \beta' - \gamma')e^{-\lambda} \ln 2 - o(1). \quad (\text{B.9})$$

If $\gamma' > k^{5/2}e^{-\lambda}$, then by Lemma B.3 and our assumption on g

$$\frac{1}{n} \ln \mathbb{E} [Z_{\beta', \gamma'}] \leq h(\beta') + O_k(k4^{-k}) - \frac{\ln k}{6} \gamma' e^{-\lambda} \leq (1 - \beta')e^{-\lambda} \ln 2 - \frac{\ln k}{6} \gamma' e^{-\lambda},$$

Thus, by applying (B.9), we infer that $Z_{\beta', \gamma'} > \exp(\Omega(n)) \mathbb{E} [Z_{\beta', \gamma'}]$. By Markov's inequality, the probability of this event is $\exp(-\Omega(n))$. On the other hand, if $\gamma' < k^{5/2}e^{-\lambda}$, then for sufficiently large k

$$\frac{1}{n} \ln \mathbb{E} [Z_{\beta', \gamma'}] \leq h(\beta') + O_k(k4^{-k}) \leq -k^3 4^{-k+1} + (1 - \beta')e^{-\lambda} \ln 2 < -k^3 4^{-k} + (1 - \beta' - \gamma')e^{-\lambda} \ln 2.$$

Thus, again by applying (B.9), we infer that also in this case $Z_{\beta', \gamma'} > \exp(\Omega(n)) \mathbb{E} [Z_{\beta', \gamma'}]$, and Markov's inequality asserts that the probability of this event is $\exp(-\Omega(n))$.

Since the probability that either case occurs is $\exp(-\Omega(n))$, we conclude that the same is true of the event " $Z'_j > 0$ ". Taking the union bound over j then completes the induction step, i.e., $Z_{\leq \beta_{i+1}} = 0$ w.h.p. \square

Finally, the upper bound on $r_{k-\text{NAE}}$ claimed in Theorem 1.1 follows directly from Proposition B.4.

C Proof of the lower bound

C.1 Outline

Let \mathbf{d}, \mathbf{D} be as in Section 5. In the extended abstract, we presented a slightly streamlined definition of "good". Technically it will be more convenient to work with the following definition. (It will emerge later that the two definitions are equivalent.) Recall that $\lambda = kr/(2^{k-1} - 1)$.

Definition 1. We call a solution $\sigma \in \{0, 1\}^V$ of $\Phi_{\mathbf{d}}$ β -good if it satisfies the following conditions.

1. σ is β -heavy and the total number of critical clauses is equal to λn .
2. No variable supports more than $3k$ clauses.
3. We have

$$\frac{1}{n} \ln \left| \left\{ \tau \in \mathcal{S}(\Phi_{\mathbf{d}}) : \text{dist}(\sigma, \tau)/n \leq \frac{1}{2} - 2^{-k/3} \right\} \right| \leq (1 - \beta) \exp(-\lambda) \ln 2 + O_k(k^{13} 4^{-k}).$$

Let \mathcal{Z}_{β} be the number of β -good solutions. As a first step, we determine the expectation of \mathcal{Z}_{β} .

Proposition C.2. Suppose that \mathbf{d} is chosen from the distribution \mathbf{D} . Then w.h.p.

$$\frac{1}{n} \ln \mathbb{E} [\mathcal{Z}_{\beta}] \geq \frac{2\rho - \ln 2}{2^k} + f(\beta) - O_k(k^{13} 4^{-k}).$$

Let us fix an assignment $\sigma \in \{0, 1\}^V$, say $\sigma = \mathbf{1}$. Moreover, let Σ be the event that σ is a β -good solution. Let $\mathcal{Z}_\beta(t)$ be the number of β -good solutions $\tau \in \mathcal{S}(\Phi_d)$ such that $\text{dist}(\sigma, \tau) = t$. Then the symmetry properties of Φ_d imply the following.

Fact C.3. *For any d we have $\mathbb{E}[\mathcal{Z}_\beta^2] = \mathbb{E}[\mathcal{Z}_\beta|\Sigma] \cdot \mathbb{E}[\mathcal{Z}_\beta]$.*

Thus, we need to compare $\mathbb{E}[\mathcal{Z}_\beta|\Sigma]$ with $\mathbb{E}[\mathcal{Z}_\beta]$. Let $\delta = 2^{-k/3}$. By the linearity of expectation and by the symmetry of $\mathcal{S}(\Phi)$ with respect to inversion, for any d

$$\begin{aligned}
\mathbb{E}[\mathcal{Z}_\beta|\Sigma] &= \sum_{t=0}^n \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \leq 2 \sum_{t=0}^{n/2} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \\
&= 2 \sum_{t \leq (\frac{1}{2}-\delta)n} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] + 2 \sum_{(\frac{1}{2}-\delta)n < t \leq \frac{1}{2}n} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \\
&\leq 2 \left| \left\{ \tau \in \mathcal{S}(\Phi_d) : \text{dist}(\sigma, \tau)/n \leq \frac{1}{2} - 2^{-k/3} \right\} \right| + 2 \sum_{(\frac{1}{2}-\delta)n < t \leq \frac{1}{2}n} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \\
&\leq 2 \exp \left[n((1-\beta) \exp(-\lambda) \ln 2 - O_k(k^{13}4^{-k})) \right] + 2 \sum_{(\frac{1}{2}-\delta)n < t \leq \frac{1}{2}n} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \quad (\text{C.1})
\end{aligned}$$

by the definition of β -good. Let

$$r^* = 2^{k-1} \ln 2 - \left(\frac{\ln 2}{2} + \frac{1}{4} \right) - k^{14}2^{-k}.$$

Lemma C.4. *For any $r < r^*$ there exists $0 < \beta \leq \frac{1}{2}$ such that for d chosen from \mathbf{D} w.h.p.*

$$\mathbb{E}[\mathcal{Z}_\beta] \geq \exp \left[n((1-\beta)e^{-\lambda} \ln 2) + k^{14}2^{-k+1} \right].$$

Proof. This follows from Proposition C.2 and a little bit of calculus. \square

As a next step, we are going to bound the second summand in (C.1). This is technically the most demanding part of this work. In Appendix D we are going to prove the following.

Lemma C.5. *Let $\delta = 2^{-k/3}$. There is a number $C = C(k)$ such that for a degree sequence d chosen from \mathbf{D} we have w.h.p.*

$$\sum_{(\frac{1}{2}-\delta)n < t \leq \frac{1}{2}n} \mathbb{E}[\mathcal{Z}_\beta(t)|\Sigma] \leq C \cdot \mathbb{E}[\mathcal{Z}_\beta].$$

Corollary C.6. *For any $r < r^*$ there is $0 < \beta \leq \frac{1}{2}$ such that $\mathbb{E}[\mathcal{Z}_\beta|\Sigma] \leq C \cdot \mathbb{E}[\mathcal{Z}_\beta]$ for some constant $C = C(k) > 1$.*

Proof. This follows directly from (C.1), Lemma C.4, and Lemma C.5. \square

Proof of Theorem 1.1 (lower bound). By Corollary C.6 and the Paley-Zygmund inequality, for any $r < r^*$ for a random d chosen from the distribution \mathbf{D} we have w.h.p.

$$\mathbb{P}[\Phi_d \text{ has an NAE-solution}] \geq \mathbb{P}[\mathcal{Z}_\beta > 0] \geq 1/C. \quad (\text{C.2})$$

Since \mathbf{D} is precisely the distribution of the degree sequence of the uniformly random formula Φ , we have

$$\mathbb{E}_{\mathbf{d}} [\mathbb{P} [\Phi_{\mathbf{d}} \text{ has an NAE-solution}]] = \mathbb{P} [\Phi \text{ has an NAE-solution}],$$

where the expectation on the left hand side ranges over \mathbf{d} chosen from \mathbf{D} . Therefore, (C.2) implies that

$$\mathbb{P} [\Phi \text{ has an NAE-solution}] \geq \frac{1}{C} - o(1), \quad (\text{C.3})$$

which remains bounded away from 0 as $n \rightarrow \infty$. Hence, (C.3) implies that $r_{k\text{-NAE}} \geq r^*$, as the k -NAESAT threshold is sharp. \square

C.2 Proof of Proposition C.2

We begin with the following simple observation.

Lemma C.7. *For any \mathbf{d} and any $\sigma \in \{0, 1\}^V$ we have $\mathbb{P} [\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})] = (1 - 2^{1-k})^m$.*

Proof. We may assume without loss that $\sigma = \mathbf{1}$. Then σ is a solution iff each clause has both a positive and a negative literal. Since the *signs* of the literals are chosen uniformly and independently, the assertion follows. \square

We defer the proof of the following result to Section C.3.

Proposition C.8. *Let \mathbf{d} be chosen from \mathbf{D} . Then w.h.p. we have*

$$\frac{1}{n} \ln \mathbb{P} [\sigma \text{ has Properties 1. and 2. from Definition 1} \mid \sigma \in \mathcal{S}(\Phi_{\mathbf{d}})] = f(\beta) + O_k(k4^{-k}).$$

To continue, we need the following basic fact about the random degree distribution \mathbf{d} . For a set $S \subset V$ we let $\text{Vol}(S) = \sum_{x \in S} d_x$.

Lemma C.9. *Let \mathbf{d} be chosen from \mathbf{D} . Then w.h.p. the following is true.*

$$\text{For any set } S \subset V \text{ we have } \text{Vol}(S) \leq 10 \max \{kr|S|, |S| \ln(n/|S|)\}. \quad (\text{C.4})$$

Proof. For any fixed $S \subset V$ the volume $\text{Vol}(S) = \sum_{x \in S} d_x$ is a sum of independent Poisson variables $\text{Po}(kr)$. Hence, $\text{Vol}(S) = \text{Po}(|S|kr)$, and the lemma follows from a straight first moment argument. \square

Let us call $S \subset V$ *dense* if each variable in S supports at least two clauses that each feature another variable from S .

Lemma C.10. *Let \mathbf{d} be chosen from \mathbf{D} and let $\sigma \in \{0, 1\}^V$. Let \mathcal{A} be the event that $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ and that σ satisfies conditions 1.–2. in Definition 1. Then w.h.p.*

$$\mathbb{P} [\text{there is a dense } S \subset V, |S| \leq n/k^5 \mid \mathcal{A}] = o(1).$$

Proof. We may assume that \mathbf{d} satisfies (C.4). Let $\mathcal{D}(S)$ be the event that $S \subset V$ is dense. We claim that

$$\mathbb{P}_{\Phi_{\mathbf{d}}} [\mathcal{D}(S)] \leq k^{2|S|} \cdot \frac{(k-1)^{2|S|} \text{Vol}(S)^{2|S|}}{(krn/2)^{2|S|}} \leq \left(\frac{2k^2 \text{Vol}(S)}{krn} \right)^{2|S|}.$$

Indeed, the factor $k^{2|S|}$ accounts for the number of ways to choose the two relevant clauses supported by each variable, and the second factor bounds the probability that each of these clauses contains another occurrence of a variable from S . Now, (C.4) yields

$$\mathbb{P}_{\Phi_d}[\mathcal{D}(S)] \leq \left(\frac{2k^2|S| \ln(n/|S|)}{n} \right)^{2|S|}.$$

For $0 < s \leq 1/k^5$ let X_s be the number of sets S of size $|S| = sn$ for which $\mathcal{D}(S)$ occurs. Then

$$\mathbb{E}[X_s] \leq \binom{n}{sn} [2k^2 s \ln(1/s)]^{2sn} \leq \left[\frac{e}{s} \cdot (2k^2 s \ln(1/s))^2 \right]^{sn} \leq (4ek^4 s \ln^2(1/s))^{sn} = o(1).$$

Summing over all possible s and using Markov's inequality completes the proof. \square

Lemma C.11. *The expected number of solutions $\sigma \in \mathcal{S}(\Phi)$ in which more than $k^4 2^{-k} n$ variables support at most four clauses is $\leq \exp(-nk^3/2^k)$.*

Proof. Fix an assignment $\sigma \in \{0, 1\}^V$, say $\sigma = \mathbf{1}$. Then number of clauses supported by each $x \in V$ is asymptotically Poisson with mean λ . Let \mathcal{E}_x be the event that x supports no more than three clauses. Then

$$\mathbb{P}[\mathcal{E}_x] \leq \lambda^3 \exp(-\lambda) \leq k^4 2^{-k-1}.$$

The events $(\mathcal{E}_x)_{x \in V}$ are negatively correlated. Therefore, the total number X of variables $x \in V$ for which \mathcal{E}_x occurs is stochastically dominated by a binomial variable $\text{Bin}(n, k^4 2^{-k-1})$. Hence, the assertion follows from Chernoff bounds. \square

Let us call a set $S \subset V$ *self-contained* if each variable in S supports at least two clauses that consist of variables in S only. There is a simple process that yields a (possibly empty) self-contained set S .

- For each variable x that supports at least one clause, choose such a clause C_x randomly.
- Let R be the set of all variables that support at least four clauses.
- While there is a variable $x \in R$ that supports fewer than two clauses $\Phi_i \neq C_x$ that consist of variables of R only, remove x from R .

The clauses C_x will play a special role later.

Lemma C.12. *The expected number of solutions $\sigma \in \mathcal{S}(\Phi)$ for which the above process yields a set R of size $|R| \leq (1 - k^5/2^k)n$ is bounded by $\exp(-\Omega(n))$.*

Proof. Let $\sigma \in \{0, 1\}^V$ be an assignment, say $\sigma = \mathbf{1}$. Let Q be the set of all variables that support fewer than three clauses. By Lemma C.11 we may condition on $|Q| \leq k^4 2^{-k} n$. Assume that its size is $|R| \leq (1 - k^5/2^k)n$. Then there exists a set $S \subset V \setminus (R \cup Q)$ of size $\frac{1}{2}k^5 n/2^k \leq S \leq k^5 n/2^k$ such that each variable in S supports two clauses that contain another variable from $S \cup Q$. With $s = |S|/n$ the probability of this event is bounded by

$$\binom{m}{2sn} \left[\frac{2^{1-k}}{1 - 2^{1-k}} \cdot \frac{k^2 |S \cup Q|^2}{n^2} \right]^{2sn} \leq [4ek^2 s]^{2sn}.$$

Hence, the expected number of set S for which the aforementioned event occurs is bounded by

$$\binom{n}{s} [4ek^2 s]^{2sn} \leq \left[\frac{e}{s} \cdot (4ek^2 s)^2 \right]^{sn} \leq \exp(-sn).$$

Since $\mathbb{E}[Z(\Phi)] \leq \exp(O_k(2^{-k}n))$, the assertion follows. \square

Corollary C.13. *Let \mathbf{d} be chosen from \mathbf{D} . Then the expected number of solutions $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ for which the above process yields a set R of size $|R| \leq (1 - k^5/2^k)n$ is bounded by $\exp(-\Omega(n))$.*

Proof. Since the random formula Φ can be generated by first choosing \mathbf{d} from \mathbf{D} and then generating $\Phi_{\mathbf{d}}$, the assertion follows from Lemma C.12. \square

Let us call a variable x is *attached* if x supports a clause whose other $k - 1$ variables belong to R .

Corollary C.14. *W.h.p. a degree sequence \mathbf{d} chosen from \mathbf{D} has the following property. Let $\sigma \in \{0, 1\}^V$ and let \mathcal{A} be the event that $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ and that σ satisfies Conditions 1. and 2. in Definition 1. Moreover, let Y be the number variables that support a clause but that are not attached. Then*

$$\mathbb{P}_{\Phi_{\mathbf{d}}} \left[Y \leq nk^{13}4^{-k} \mid \mathcal{A} \right] = 1 - o(1).$$

Proof. We may assume that \mathbf{d} satisfies (C.4). Let $F = V \setminus R$. Then (C.4) ensures that $\frac{\text{Vol}(F)}{k^n} \leq \frac{2k^6}{2^k}$. Therefore, for each of the “special” clause \mathcal{C}_x that we reserved for each x that supports at least one clause the probability of containing a variable from $F \setminus \{x\}$ is bounded by

$$(1 + o_k(1))k \cdot \frac{\text{Vol}(F)}{k^n} \leq \frac{3k^7}{2^k}.$$

Furthermore, these events are negatively correlated (due to the bound on $\text{Vol}(F)$). Since $|V \setminus R| \leq k^5 n / 2^k$ w.h.p. by Corollary C.13, the assertion thus follows from Chernoff bounds. \square

Let us call a variable $x \in V$ ξ -*rigid* in a solution $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ if for any solution $\tau \in \mathcal{S}(\Phi_{\mathbf{d}})$ with $\tau(x) \neq \sigma(x)$ we have $\text{dist}(\sigma, \tau) \geq \xi n$.

Corollary C.15. *W.h.p. a degree sequence \mathbf{d} chosen from \mathbf{D} has the following property. Let $\sigma \in \{0, 1\}^V$ and let \mathcal{A} be the event that $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ and that σ satisfies Conditions 1. and 2. in Definition 1. Moreover, let Y be the number of variables that support a clause but that are k^{-5} -rigid. Then*

$$\mathbb{P}_{\Phi_{\mathbf{d}}} \left[Y \leq nk^{13}4^{-k} \mid \mathcal{A} \right] = 1 - o(1).$$

Proof. We condition on the event \mathcal{A} . By Corollary C.13, we may assume that the self-contained set R has size $|R| \geq (1 - k^5/2^k)n$. Assume that there is $\tau \in \mathcal{S}(\Phi_{\mathbf{d}})$, $\text{dist}(\sigma, \tau) < n/k^5$, such that

$$\Delta = \{x \in R : \tau(x) \neq \sigma(x)\}$$

is non-empty. Then Δ is dense. Indeed, every $x \in \Delta$ supports at least two clauses, and thus Δ must contain another variable from each of them. Thus, Lemma C.10 shows that $|\Delta| \geq n/k^5$, which is a contradiction.

Hence, w.h.p. all variables $x \in R$ are k^{-5} -rigid. Furthermore, if a variable y is attached, then for any solution τ with $\tau(y) \neq \sigma(y)$ there is $x \in R$ such that $\tau(x) \neq \sigma(x)$. Consequently, all attached variables are k^{-5} -rigid w.h.p. Therefore, the assertion follows from Corollary C.14. \square

To complete the proof, we need the following fairly simple lemma.

Lemma C.16. *The expected number of pairs of solutions $\sigma, \tau \in \mathcal{S}(\Phi)$ such that $\frac{n}{k^6} \leq \text{dist}(\sigma, \tau) \leq (\frac{1}{2} - 2^{-k/2})n$ is $\leq \exp(-\Omega(n))$.*

Proof. For a given $0 \leq \alpha \leq 1$ let P_α denote the number of pairs $\sigma, \tau \in \mathcal{S}(\Phi)$ with $\text{dist}(\sigma, \tau) = \alpha n$. As worked out in [3], we have

$$\frac{1}{n} \ln \mathbb{E}[P_\alpha] \leq \ln 2 - \alpha \ln \alpha - (1 - \alpha) \ln(1 - \alpha) + r \ln \left(1 - 2^{2-k} + 2^{1-k}(\alpha^k + (1 - \alpha)^k) \right).$$

It is a mere exercise in calculus to verify that the r.h.s. is strictly negative for all $k^{-6} \leq \alpha \leq \frac{1}{2} - 2^{-k/2}$. \square

Corollary C.17. *W.h.p. a degree sequence \mathbf{d} chosen from \mathbf{D} has the following property. The expected number of pairs of solutions $\sigma, \tau \in \mathcal{S}(\Phi_{\mathbf{d}})$ such that $\frac{n}{k^6} \leq \text{dist}(\sigma, \tau) \leq (\frac{1}{2} - 2^{-k/2})n$ is $\leq \exp(-\Omega(n))$.*

Combining Lemma C.7, Proposition C.8, Corollary C.15, and Corollary C.17, we obtain

Corollary C.18. *W.h.p. a degree sequence \mathbf{d} chosen from \mathbf{D} has the following property. Let $\sigma \in \{0, 1\}^V$ and let \mathcal{A} be the event that $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$ and that σ satisfies Conditions 1. and 2. in Definition 1. Then*

$$\mathbb{P}_{\Phi_{\mathbf{d}}} [3. \text{ in Definition 1 is satisfied} \mid \mathcal{A}] = 1 - o(1).$$

Finally, Proposition C.2 is a direct consequence of Lemma C.7, Proposition C.8, and Corollary C.18.

C.3 Proof of Proposition C.8

Let us begin with establishing the probable properties of \mathbf{d} that we will need.

Lemma C.19. *Let $\mathbf{d} = (d_1, \dots, d_n)$ be from the distribution $\mathbf{D} = \mathbf{D}(k, r, n)$. Then, with high probability, for any $0 \leq \alpha \leq (kr)^{1/2}$, the sequence \mathbf{d} has the following properties. First, for all i such that $|i - kr| \leq \alpha\sqrt{kr}$*

$$D_i = |\{j : d_j = i\}| = (1 + o(1)) \Pr[\text{Po}(kr) = i] n. \quad (\text{C.5})$$

Moreover, the remaining variables satisfy

$$D^{\geq \alpha} = \left| \left\{ j : |d_j - kr| \geq \alpha\sqrt{kr} \right\} \right| \leq 2e^{-\alpha^2/2} n \quad \text{and} \quad \sum_{j \in D^{\geq \alpha}} d_j \leq 2e^{-\alpha^2/2} (kr)^2 n. \quad (\text{C.6})$$

Proof. Let P_1, \dots, P_n be independent $\text{Po}(kr)$ random variables, and note that the joint distribution of (d_1, \dots, d_n) and (P_1, \dots, P_n) , conditional on $\sum_{1 \leq i \leq n} P_i = krn$, coincide. Since the expectation of the sum of the P_i 's equals krn , Lemma A.1 applied with $\delta = 0$ implies that for any event \mathcal{E} we have that

$$\Pr[\mathbf{d} \in \mathcal{E}] = \Pr \left[(P_1, \dots, P_n) \in \mathcal{E} \mid \sum_{1 \leq i \leq n} P_i = krn \right] = O(n^{1/2}) \Pr[(P_1, \dots, P_n) \in \mathcal{E}].$$

In other words, it is sufficient to show that the statements in the lemma hold with probability $1 - o(n^{-1/2})$ for a sequence of independent Poisson random variables. The statements follow from the Chernoff bounds and the fact that for any $\lambda = kr$ and α as assumed

$$\Pr[\text{Po}(\lambda) \geq \alpha\sqrt{\lambda}] \leq 2e^{-\alpha^2} \quad \text{and} \quad \sum_{j: |j-\lambda| \geq \alpha\sqrt{\lambda}} j \Pr[\text{Po}(\lambda) = j] \leq 2e^{-\alpha^2/2} \lambda^2.$$

\square

The aim of this section is to show that for any \mathbf{d} satisfying the conclusions of Lemma C.19

$$\frac{1}{n} \ln \Pr[\sigma \text{ has Properties 1. and 2. from Definition 1} \mid \sigma \in \mathcal{S}(\Phi_{\mathbf{d}})] = f(\beta) + O_k(4^{-k}), \quad (\text{C.7})$$

i.e., Proposition C.8 holds. We will assume that $\sigma = 1$ throughout.

First of all, let C denote the number of critical clauses. Given that $\mathbf{1}$ is a NAE-satisfying assignment, then there are for each clause in total $2^k - 2$ ways to choose the signs of the variables, each one of them being equally likely. Since the number of ways to choose the signs so as to obtain a critical clause is $2k$, the probability that a given clause is critical is $k/(2^{k-1} - 1)$. Moreover, the events that different clauses are critical are independent, implying that C is distributed like $\text{Bin}(m, k/(2^{k-1} - 1))$.

Note that $\mathbb{E}[C \mid \mathbf{1} \in \mathcal{S}(\Phi_{\mathbf{d}})] = m \cdot k/(2^{k-1} - 1) = \lambda n$. By applying Lemma A.1 with $\delta = 0$ we thus obtain that

$$\Pr[C = \lambda n \mid \mathbf{1} \in \mathcal{S}(\Phi_{\mathbf{d}})] = \Theta(n^{-1/2}).$$

It follows that the probability in (C.7) equals

$$\Theta(n^{-1/2}) \cdot \Pr[\mathbf{1} \text{ is } \beta\text{-heavy and no variable supports } \geq 3k \text{ clauses} \mid C = \lambda n \text{ and } \mathbf{1} \in \mathcal{S}(\Phi_{\mathbf{d}})]. \quad (\text{C.8})$$

In the sequel we adopt a different formulation of this probabilistic question that is based on the classical occupancy problem. Let us think of the variables as bins, such that the i th bin has capacity d_i , where $\mathbf{d} = (d_1, \dots, d_n)$. In other words, we assume that the i th bin contains d_i distinguished “slots”. Then we throw randomly λn balls into the bins, i.e., the j th ball chooses uniformly at random one of the remaining $\sum_{1 \leq i \leq n} d_i - (j-1) = krn - j + 1$ available slots, for each $1 \leq j \leq \lambda n$. In this setting, the probability in (C.8) is equal to the probability that in the balls-into-bins game with the given capacity constraints the number of empty bins equals $(1 - \beta)e^{-\lambda n}$, and no bin contains more than $2k$ balls. More precisely, let R_i , where $1 \leq i \leq n$, denote the number of balls selected from the i th bin. Then, the probability in (C.8) equals

$$\Pr[\mathcal{A} \text{ and } \mathcal{B}], \quad \text{where } \mathcal{A} = “|\{i : R_i = 0\}| = (1 - \beta)e^{-\lambda n}” \text{ and } \mathcal{B} = “\forall 1 \leq i \leq n : R_i \leq 3k”.$$

We will show that the probability above is $\exp\{(f(\beta) + O_k(4^{-k}))n\}$, which together with (C.8) completes the proof of (C.7).

In order to compute the probability of the event “ \mathcal{A} and \mathcal{B} ” we resort to the following experiment. Instead of throwing λn balls into the available slots, we decide for each slot *independently* with probability λ/kr whether it receives a ball or not. Let T be the total number of balls that are thrown in this setting, and let $B_i \sim \text{Bin}(d_i, \lambda/kr)$ be the number of balls that the i th bin received. Since the total number of slots is krn , we have that $\mathbb{E}[T] = \lambda n$. Moreover, conditional on any value of T , the T slots that receive a ball are a random subset of size T of all available slots. Thus, conditional on “ $T = \lambda n$ ” the joint distributions of (R_1, \dots, R_n) and (B_1, \dots, B_n) coincide, and by abbreviating $X_i = |\{j : B_j = i\}|$ we obtain that

$$\Pr[\mathcal{A} \text{ and } \mathcal{B}] = \Pr[X_0 = (1 - \beta)e^{-\lambda n} \text{ and } X_{>3k} = 0 \mid T = \lambda n]. \quad (\text{C.9})$$

Before we estimate the latter probability, let us give some intuitive explanation why this should be equal to $e^{(f(\beta) + O_k(4^{-k}))n}$, i.e., why the conclusion of the proposition is true. Our assumption on the bin capacities (C.5) guarantees that most bins have a capacity very close to $kr \approx k2^{k-1} \ln 2$. Recall also that the probability that any slot receives a ball is $\lambda/kr \approx 2^{-k+1}$. This means that the expected number of balls that a typical bin receives is $\approx k$, which is far smaller than the capacity of that bin. But we can say even more: since the number of balls that are received by a typical bin is $\approx \text{Bin}(kr, \lambda/kr)$, and the expected value is far less than kr , it is reasonable to assume that this number can be approximated well by a $\text{Po}(\lambda)$ distribution. So, the probability that a bin remains empty is close to $e^{-\lambda}$, and then the probability that the number of

empty bins is exactly $(1 - \beta)e^{-\lambda}n$ should be close to $\Pr[\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n]$. The argument then completes by applying Lemma B.2.

Let us now put the above intuitive reasoning on a rigorous ground. First of all, note that in the right-hand side of (C.9) the condition “ $T = \lambda n$ ” is *global*, in the sense that it binds the values of all variables B_1, \dots, B_n . We can get rid of this global restriction by applying the law of total probability. We obtain that

$$\begin{aligned} \Pr[\mathcal{A} \text{ and } \mathcal{B}] &= \frac{\Pr[T = \lambda n \text{ and } X_0 = (1 - \beta)e^{-\lambda}n \text{ and } X_{>3k} = 0]}{\Pr[T = \lambda n]} \\ &= \Pr\left[T = \lambda n \text{ and } X_0 = (1 - \beta)e^{-\lambda}n \mid X_{>3k} = 0\right] \frac{\Pr[X_{>3k} = 0]}{\Pr[T = \lambda n]}. \end{aligned} \quad (\text{C.10})$$

The remainder of the proof is devoted to showing the following bounds.

$$\Pr[T = \lambda n] = \Theta(n^{-1/2}), \quad (\text{C.11})$$

$$\Pr[X_{>3k} = 0] \geq e^{-O_k(4^{-k})n}, \quad (\text{C.12})$$

$$\Pr\left[T = \lambda n \text{ and } X_0 = (1 - \beta)e^{-\lambda}n \mid X_{>3k} = 0\right] \geq \Pr[\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n] \cdot e^{-O_k(k4^{-k})n}. \quad (\text{C.13})$$

The three inequalities together with (C.10) imply that

$$\Pr[\mathcal{A} \text{ and } \mathcal{B}] \geq \Pr[\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n] \cdot e^{-O_k(k4^{-k})n},$$

and the proof of the proposition is completed after applying Lemma B.2.

In the remainder of the proof we will write \mathcal{D}_i for the set of bins with capacity i and $\mathcal{D}^{\geq \alpha}$ for the set of bins with capacity smaller than $kr - \alpha\sqrt{kr}$ or larger than $kr + \alpha\sqrt{kr}$, and note that $|\mathcal{D}_i| = D_i$ and $|\mathcal{D}^{\geq \alpha}| = D^{\geq \alpha}$.

Proof of (C.11). Since T is distributed like $\text{Bin}(krn, \lambda/kr)$ we have that $\mathbb{E}[T] = \lambda n$. The result then follows by applying Lemma A.1 with $\delta = 0$ to T .

Proof of (C.12). Recall that the number of bins with capacity i is denoted by D_i . Since the number of balls in a bin with capacity i is distributed like $\text{Bin}(i, \lambda/kr)$, and these variables are all independent, we obtain that

$$\begin{aligned} \Pr[X_{>3k} = 0] &= \prod_{i \geq 0} \Pr[\text{Bin}(i, \lambda/kr) \leq 3k]^{D_i} \\ &\geq \prod_{i: |i - kr| < k\sqrt{kr}} \Pr[\text{Bin}(i, \lambda/kr) = 0]^{D_i} \cdot \prod_{i: |i - kr| \geq k\sqrt{kr}} \Pr[\text{Bin}(i, \lambda/kr) \leq 3k]^{D_i}. \end{aligned} \quad (\text{C.14})$$

Our assumption (C.6) guarantees that \mathbf{d} is such that

$$\sum_{i: |i - kr| \geq k\sqrt{kr}} iD_i = \sum_{j \in \mathcal{D}^{\geq k}} d_j \leq 2e^{-k^2/2}(kr)^2n.$$

Thus, if k is sufficiently large, the last term in (C.14) can be bounded with

$$\prod_{i: |i - kr| \geq k\sqrt{kr}} \Pr[\text{Bin}(i, \lambda/kr) = 0]^{D_i} = \prod_{j \in \mathcal{D}^{\geq k}} \left(1 - \frac{\lambda}{kr}\right)^{d_j} = \left(1 - \frac{\lambda}{kr}\right)^{\sum_{j \in \mathcal{D}^{\geq k}} d_j} \geq e^{-e^{-k^2/3}n}. \quad (\text{C.15})$$

Let us now consider the terms involving all i such that $|i - kr| < k\sqrt{kr}$ in (C.15). By using the estimate $\binom{a}{b} \leq (ea/b)^b$ we infer that for any such i and sufficiently large k we have

$$\Pr \left[\text{Bin} \left(i, \frac{\lambda}{kr} \right) > 3k \right] \leq \binom{i}{3k} \left(\frac{\lambda}{kr} \right)^{3k} \leq \left(\frac{ei}{3k} \frac{\lambda}{kr} \right)^{3k} \leq \left(\frac{ekr(1+o_k(1))}{3k} \frac{k \ln 2(1+o_k(1))}{kr} \right)^{3k} \leq 4^{-k}. \quad (\text{C.16})$$

Thus, since $\sum_{i \geq 0} D_i = n$, by using the fact $1 - x = e^{-x - \Theta(x^2)}$, valid for all $|x| \leq 1$,

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr [\text{Bin}(i, \lambda/kr) \leq 3k]^{D_i} \geq \prod_{i: |i-kr| < k\sqrt{kr}} (1 - 4^{-k})^{D_i} = e^{-4^{-k}n - \Theta_k(4^{-2k})n}.$$

This result, together with (C.15) and (C.14) finally prove (C.12).

Proof of (C.13). Note that

$$\Pr[\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda}n] = \binom{n}{(1 - \beta)e^{-\lambda}n} (e^{-\lambda})^{(1 - \beta)e^{-\lambda}n} (1 - e^{-\lambda})^{(1 - (1 - \beta)e^{-\lambda})n}. \quad (\text{C.17})$$

In the following proof we will approximate the probability of the event “ $T = \lambda n$ and $X_0 = (1 - \beta)e^{-\lambda}n$ ”, conditional on $X_{>3k} = 0$, by the right-hand side of the above equation times an error term, which is of order $\exp\{-O_k(k4^{-k})n\}$. In particular, we will identify the most relevant objects that contribute precisely these terms to the desired probability.

In order to prove a lower bound for the probability of the event “ $T = \lambda n$ and $X_0 = (1 - \beta)e^{-\lambda}n$ ” we will consider only specific configurations of balls that lead to the desired outcome. More precisely, let $\mathbf{b} = (b_1, \dots, b_n)$ denote a possible outcome of the random experiment that we study, where b_i denotes the number of balls in the i th bin. We will call \mathbf{b} *balanced* if it has the following properties:

1. Let $j \in \mathcal{D}_i$, where $|i - kr| \geq k\sqrt{kr}$. Then $b_j = 0$. Informally, the $D^{\geq k}$ bins with “too small” or “too big” capacities are empty.
2. Let \mathcal{D}'_i denote the set of bins in \mathcal{D}_i that do not receive a ball. For all i such that $|i - kr| < k\sqrt{kr}$

$$D'_i = |\mathcal{D}'_i| = \frac{D_i((1 - \beta)e^{-\lambda} - D^{\geq k}/n)}{1 - D^{\geq k}/n}.$$

Informally, the fraction of empty bins among those in \mathcal{D}_i is the same (and approximately equal to $(1 - \beta)e^{-\lambda}$) for all relevant i .

3. Let T_i denote the total number of balls in all bins in \mathcal{D}_i . Then, for all i such that $|i - kr| < k\sqrt{kr}$

$$T_i = t_i = \frac{D_i - D'_i}{1 - (1 - \beta)e^{-\lambda}} \frac{\lambda i}{kr} \cdot x,$$

where x is chosen such that the sum of all t_i is λn . As we shall see later, see (C.27), x is very close to 1. Then again, informally this requires that the fraction of balls in the bins in \mathcal{D}_i is approximately λ for all relevant i .

4. For all $1 \leq i \leq n$ we have $b_i \leq 3k$, i.e., $X_{>3k}(\mathbf{b}) = 0$.

By our construction, note that if \mathbf{b} is balanced, then $X_0(\mathbf{b}) = (1 - \beta)e^{-\lambda}n$ and $T(\mathbf{b}) = \lambda n$. Thus,

$$\Pr[T = \lambda n \text{ and } X_0 = (1 - \beta)e^{-\lambda}n \mid X_{>3k} = 0] \geq \Pr[(B_1, \dots, B_n) \text{ is balanced}]. \quad (\text{C.18})$$

In the sequel we will estimate the latter probability. First of all, note that the number of ways to choose the empty bins in a balanced \mathbf{b} is

$$\prod_{i: |i-kr| < k\sqrt{kr}} \binom{D_i}{D'_i}. \quad (\text{C.19})$$

Note that bins contained in $\mathcal{D}^{\geq k}$ do not have to be counted explicitly, since they are contained in the set of empty bins per definition. Let us write $\text{Bin}_{i,j}(N, p)$ for a binomially distributed random variable that is conditioned on being in the interval $[i, j]$. Then, after having fixed the locations of the empty bins, the probability that (B_1, \dots, B_n) is balanced with precisely the chosen set of empty bins is

$$\prod_{i: |i-kr| \geq k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) = 0 \right]^{D_i} \cdot \prod_{i: |i-kr| < k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) = 0 \right]^{D'_i} \Pr [\mathcal{T}_i \mid X_{>3k} = 0], \quad (\text{C.20})$$

where \mathcal{T}_i is the event “ $T_i = t_i$ and $\forall j \in \mathcal{D} \setminus \mathcal{D}'_i : B_j \geq 1$ ”. Let T'_i be a sum of $D_i - D'_i$ independent variables, which are distributed like $\text{Bin}_{1,3k}(i, \lambda/rk)$. Then

$$\Pr[\mathcal{T}_i \mid X_{>3k} = 0] = \Pr \left[T'_i = \frac{D_i - D'_i}{1 - (1 - \beta)e^{-\lambda}} \frac{\lambda i}{kr} \cdot x \right] \Pr[\text{Bin}_{0,3k}(i, \lambda/rk) \geq 1]^{D_i - D'_i}. \quad (\text{C.21})$$

The probability that (B_1, \dots, B_n) is balanced is then the product of the terms in (C.19) and (C.20). In the remaining proof we will estimate the five terms in (C.19)–(C.21).

We begin with estimating the product in (C.19). Let α be such that $D'_i = \alpha D_i$, and note that α is independent of i . Since $0 \leq D^{\geq k} \leq 2e^{-k^2/2}n$, see (C.6), we obtain that

$$\alpha = \frac{(1 - \beta)e^{-\lambda} - D^{\geq k}/n}{1 - D^{\geq k}/n} = (1 - \beta)e^{-\lambda} + \Theta(1)e^{-k^2/2}. \quad (\text{C.22})$$

By applying Proposition A.2 with $\alpha = (1 - \beta)e^{-\lambda}$ and $\varepsilon = \Theta(1)e^{-k^2/2}$ we infer that

$$\prod_{i: |i-kr| < k\sqrt{kr}} \binom{D_i}{D'_i} = \prod_{i: |i-kr| < k\sqrt{kr}} \frac{\Theta(1)}{\sqrt{\alpha(1 - \alpha)D_i}} e^{(H(\alpha) + O_k(ke^{-k^2/2}))D_i} = e^{H(\alpha)(n - D^{\geq k}) + O_k(ke^{-k^2/2})n}.$$

By using once more the fact $0 \leq D^{\geq k} \leq 2e^{-k^2/2}n$ and by applying Proposition A.2 we infer that

$$\prod_{i: |i-kr| < k\sqrt{kr}} \binom{D_i}{D'_i} = \binom{n}{(1 - \beta)e^{-\lambda}n} \cdot e^{O_k(e^{-k^2/3})n}. \quad (\text{C.23})$$

This estimate contributes the binomial coefficient in (C.17) to our lower bound for the probability in (C.18). It remains to bound the expression in (C.20). Let us begin with considering the first product, which accounts for all i that deviate significantly from kr . Since $\Pr[\text{Bin}_{i,j}(N, p) = \ell] \geq \Pr[\text{Bin}(N, p) = \ell]$ for all N, p, i, j and $i \leq \ell \leq j$ we have

$$\prod_{i: |i-kr| \geq k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) = 0 \right]^{D_i} \stackrel{(\text{C.15})}{\geq} e^{-e^{-k^2/3}n}. \quad (\text{C.24})$$

Let us consider the middle term in (C.20). Using again the property $\Pr[\text{Bin}_{i,j}(N, p) = \ell] \geq \Pr[\text{Bin}(N, p) = \ell]$ and the facts $1 - x = e^{-x - \Theta(x^2)}$ and $\lambda = k \ln 2 + O_k(k2^{-k})$ and $r = 2^{k-1} \ln 2 - c$ we obtain

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) = 0 \right]^{D'_i} \geq \exp \left\{ - \left(\frac{\lambda}{kr} + O_k(4^{-k}) \right) \alpha \sum_{i: |i-kr| < k\sqrt{kr}} i D_i \right\}.$$

By using again the property of \mathbf{d} in (C.6) we infer that

$$\sum_{i: |i-kr| < k\sqrt{kr}} iD'_i = krn - \sum_{j \in \mathcal{D}^{\geq k}} d_j = krn - O_k(e^{-k^2/2})n.$$

Recall that $\alpha = (1 - \beta)e^{-\lambda} + \Theta(1)e^{-k^2/2}$. Thus the middle term in (C.20) is at least

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) = 0 \right]^{D'_i} \geq (e^{-\lambda})^{(1-\beta)e^{-\lambda}n} \cdot e^{-O_k(k4^{-k})}. \quad (\text{C.25})$$

This estimate contributes the $(e^{-\lambda})^{(1-\beta)e^{-\lambda}n}$ term in (C.17) to our lower bound for the probability in (C.18). We finally consider the probability of the event \mathcal{T}_i in (C.20), c.f. also (C.21). The last term in (C.21) can be bounded as follows. First, note that

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr[\text{Bin}_{0,3k}(i, \lambda/rk) \geq 1]^{D_i - D'_i} \geq \prod_{i: |i-kr| < k\sqrt{kr}} \Pr[1 \leq \text{Bin}(i, \lambda/rk) \leq 3k]^{D_i - D'_i}$$

By using (C.16) and the fact $1 - x = e^{-x - \Theta(x^2)}$, where $0 \leq x \leq 1$, we obtain

$$\Pr[1 \leq \text{Bin}(i, \lambda/rk) \leq 3k] \geq 1 - (1 - \lambda/rk)^i - 4^{-k} = \exp\{-(1 - \lambda/rk)^i + O_k(4^{-k})\}.$$

With this estimate at hand we can bound the last term in (C.21). We get that

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) \geq 1 \right]^{D_i - D'_i} \geq \exp \left\{ -(1 - \alpha) \sum_{i: |i-kr| < k\sqrt{kr}} (1 - \lambda/rk)^i D_i + O_k(4^{-k})n \right\}.$$

Our assumption (C.5) on \mathbf{d} guarantees that $D_i = (1 + o(1)) \Pr[\text{Po}(kr) = i]n$. Thus, the sum in the previous equation is at most

$$(1 + o(1))n \sum_{i \geq 0} (1 - \lambda/rk)^i \Pr[\text{Po}(kr) = i] = (1 + o(1))e^{-\lambda}n,$$

from which we get that, by applying again the fact $1 - x = e^{-x - \Theta(x^2)}$,

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr \left[\text{Bin}_{0,3k} \left(i, \frac{\lambda}{kr} \right) \geq 1 \right]^{D_i - D'_i} \geq (1 - e^{-\lambda})^{(1-(1-\beta)e^{-\lambda})n} \cdot e^{-O_k(k4^{-k})}. \quad (\text{C.26})$$

This estimate contributes the last missing term in (C.17) to our lower bound for the probability in (C.18).

It remains to bound the probability for the event “ $T'_i = t_i$ ” in (C.21), for all i with the property $|i - kr| < k\sqrt{kr}$. Recall that $t_i = \frac{D_i - D'_i}{1 - (1 - \beta)e^{-\lambda}} \frac{\lambda_i}{kr} \cdot x$, where x is such that the sum of the t_i 's is λn . Let us begin with estimating the value of x . Note that

$$\lambda n = x \sum_{i: |i-kr| < k\sqrt{kr}} t_i = \frac{x\lambda(1 - \alpha)}{kr(1 - (1 - \beta)e^{-\lambda})} \sum_{i: |i-kr| < k\sqrt{kr}} iD_i.$$

Recall (C.22), which guarantees that $\alpha = (1 - \beta)e^{-\lambda} + \Theta(1)e^{-k^2/2}$. Moreover, the property (C.6) allows us to assume for large k that $\sum_{j \in \mathcal{D}^{\geq k}} d_j \leq e^{-k^2/3}n$. Thus, the above equation simplifies to

$$\lambda n = \frac{x\lambda(1 - (1 - \beta)e^{-\lambda} + O_k(e^{-k^2/2}))}{kr(1 - (1 - \beta)e^{-\lambda})} (1 - O_k(e^{-k^2/3}))krn \implies x = 1 + O_k(e^{-k^2/3}). \quad (\text{C.27})$$

Let us now return to our original goal of estimating the probability for the event “ $T'_i = t_i$ ” in (C.21). Recall that T'_i is the sum of $D_i - D'_i$ independent variables, all distributed like $\text{Bin}_{1,3k}(i, \lambda/kr)$. We will apply Lemma A.1. First of all, note that

$$\mathbb{E}[\text{Bin}_{1,3k}(i, \lambda/kr)] = \frac{\frac{i\lambda}{kr} - \sum_{j>3k} j \Pr[\text{Bin}(i, \lambda/kr) = j]}{\Pr[1 \leq \text{Bin}(i, \lambda/kr) \leq 3k]} = \frac{i\lambda}{kr} + \Theta_k(k2^{-k}),$$

and similarly, since $i = \Theta(1)kr$, that

$$\sigma^2 = \text{Var}[\text{Bin}_{1,3k}(i, \lambda/kr)] = \Theta(1) \frac{i\lambda}{kr} = \Theta(\lambda).$$

Thus, the event “ $T'_i = t_i$ ” is equivalent to “ $T'_i = (D_i - D'_i)(\mathbb{E}[\text{Bin}_{1,3k}(i, \lambda/kr)] + \Theta_k(k^{1/2}2^{-k})\sigma)$ ”. By applying Lemma A.1 we arrive at

$$\prod_{i: |i-kr| < k\sqrt{kr}} \Pr[T'_i = t_i] = \exp \left\{ \sum_{i: |i-kr| < k\sqrt{kr}} (-\delta^2/2 + O(c\delta^3))(D_i - D'_i) \right\} = \exp\{-O_k(k4^{-k})n\}.$$

Combining this result with Equations (C.18)–(C.21) and (C.23)–(C.26) yields (C.13), as desired.

D Proof of Lemma C.5

D.1 Outline

Let $\sigma = \mathbf{1}$ be the all-true assignment and let \mathbf{d} be a degree sequence chosen from the distribution \mathbf{D} . Let Σ be the event that σ is a β -good solution. Furthermore, let Σ' be the event that σ is a solution that satisfies conditions 1. and 2. in Definition 1.

Fact D.1. *Let \mathbf{d} be a degree sequence chosen from the distribution \mathbf{D} . Then $\mathbb{P}[\Sigma] \sim \mathbb{P}[\Sigma']$ w.h.p.*

Proof. This is a direct consequence of Corollary C.18. □

Let $\mathcal{Z}'_\beta(t)$ be the number of solutions τ such that $\text{dist}(\sigma, \tau) = t$ that satisfy conditions 1. and 2. in Definition 1. Moreover, let \mathcal{Z}'_β be the number of all solutions τ that satisfy conditions 1. and 2. in Definition 1. For $0 \leq t \leq n/2$ we let

$$\mu(t) = \mathbb{E}[\mathcal{Z}'_\beta(t) \mid \Sigma'].$$

The main step of the proof lies in establishing the following proposition.

Proposition D.2. *There is a constant $c = c(k) > 0$ such that for \mathbf{d} chosen from \mathbf{D} the following two statements hold w.h.p.*

1. We have $\mu(n/2) \leq \frac{c}{\sqrt{n}} \cdot \mathbb{E}[\mathcal{Z}'_\beta]$.
2. For any $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ we have $\mu(\alpha n) \leq \exp\left[-c\left(\alpha - \frac{1}{2}\right)^2 n\right] \mu(n/2)$.

Proof of Lemma C.5 (assuming Proposition D.2). By Fact D.1 we have w.h.p.

$$\begin{aligned}
\sum_{(\frac{1}{2}-2^{-k/3})n \leq t \leq n/2} \mathbb{E} [\mathcal{Z}_\beta(t) | \Sigma] &\sim \sum_{(\frac{1}{2}-2^{-k/3})n \leq t \leq n/2} \mathbb{E} [\mathcal{Z}_\beta(t) | \Sigma'] \\
&\leq \sum_{(\frac{1}{2}-2^{-k/3})n \leq t \leq n/2} \mathbb{E} [\mathcal{Z}'_\beta(t) | \Sigma'] \\
&= \sum_{(\frac{1}{2}-2^{-k/3})n \leq t \leq n/2} \mu(t) \\
&\leq c' \sqrt{n} \cdot \mu(n/2) \quad [\text{by Proposition D.2, part 2, with } c' = c'(k) > 1] \\
&\leq cc' \cdot \mathbb{E} [\mathcal{Z}'_\beta] \quad [\text{by Proposition D.2, part 1}] \\
&\leq (1 + o(1))cc' \mathbb{E} [\mathcal{Z}_\beta] \quad [\text{by Fact D.1}],
\end{aligned}$$

as desired. \square

The following subsections are devoted to the proof of Proposition D.2.

D.2 The probabilistic framework

Recall that we denote the clauses of a k -CNF formula Φ by Φ_1, \dots, Φ_m , i.e., $\Phi = \Phi_1 \wedge \dots \wedge \Phi_m$. Furthermore, for each clause Φ_i we let $\Phi_{i1}, \dots, \Phi_{ik}$ signify the literals that the clause consists of, i.e., $\Phi_i = \Phi_{i1} \vee \dots \vee \Phi_{ik}$.

We are going to break down $\mu(t)$ into a sum of different terms of various types. This requires a few definitions and a bit of notation. Given the sequence $\mathbf{d} = (d_x)_{x \in V}$ chosen from the distribution \mathbf{D} , we let

$$B = \bigcup_{x \in V} \{x\} \times [d_x],$$

where $[d_v] = \{1, 2, \dots, d_v\}$. We think of the elements of B as “balls”, so that B contains d_x balls (x, j) , $j \in [d_x]$, associated with each variable x . A *configuration* is a bijection $\pi : B \rightarrow [m] \times [k]$. Furthermore, a *signature* is a map $s : [m] \times [k] \rightarrow \{\pm 1\}$.

A configuration π and a signature s give rise to a formula $\Phi(\pi, s)$ as follows: for each $(i, j) \in [m] \times [k]$

- $\Phi(s, \pi)_{ij}$ is a positive literal if $s(i, j) = 1$ and a negative literal if $s(i, j) = -1$,
- the variable underlying $\Phi(s, \pi)_{ij}$ is the variable x such that $(i, j) \in \pi(x, [d_x])$.

We let π denote a configuration chosen uniformly at random, and we let s denote a signature chosen uniformly at random and independently of π .

Fact D.3. For any event \mathcal{E} we have $\mathbb{P} [\Phi_{\mathbf{d}} \in \mathcal{E}] = \mathbb{P} [\Phi(\pi, s) \in \mathcal{E}]$.

Proof. For each formula Φ with degree sequence \mathbf{d} there are precisely $\prod_{x \in V} d_x!$ pairs (s, π) such that $\Phi = \Phi(s, \pi)$. \square

Thus, from now on we may work with the random formula $\Phi(\pi, s)$ that emerges from choosing a random configuration and independently a signature. This will be useful because some properties depend only on the signature, and thus we will be able to treat them independently of the choice of the configuration.

Let $g : B \rightarrow \{\text{red}, \text{blue}\}$ be a map that assigns a color to each ball. For each variable x we let

$$\text{red}_x(g) = |\{j \in [d_x] : g(x, j) = \text{red}\}|, \quad \text{blue}_x(g) = |\{j \in [d_x] : g(x, j) = \text{blue}\}|.$$

Furthermore, for a pair (g_σ, g_τ) of maps $B \rightarrow \{\text{red}, \text{blue}\}$ and $\tau \in \{0, 1\}^V$ we say that (σ, τ) is (g_σ, g_τ) -*valid* for a formula Φ if the following conditions are satisfied.

- Under σ each variable x supports precisely $\text{red}_x(g_\sigma)$ clauses.
- Under τ each variable x supports precisely $\text{red}_x(g_\tau)$ clauses.
- The number of clauses that any x supports under *both* σ, τ is $|\{j \in [d_x] : g_\sigma(x, j) = g_\tau(x, j) = \text{red}\}|$.

Let s be a signature and let π be a configuration. We call an assignment $\tau \in \{0, 1\}^V$ *g-valid for (s, π)* if the following two conditions are satisfied.

- $\tau \in \mathcal{S}(\Phi(s, \pi))$.
- For any $(i, j) \in [m] \times [k]$ the following is true. Let $(u, v) = \pi(i, j)$. Then $g(i, j) = \text{red}$ iff $|\Phi(s, \pi)_{uv}|$ supports $|\Phi(s, \pi)_u|$.

In words, τ is *g-valid for (s, π)* if τ is a solution of the formula $\Phi(s, \pi)$ induced by s, π , and if each ball (i, j) that is colored red under g supports the clause that it is mapped to under π , and vice versa.

Fact D.4. Let $g_\sigma, g_\tau : B \rightarrow \{\text{blue}, \text{red}\}$. Then

$$\mathbb{P}[(\sigma, \tau) \text{ is } (g_\sigma, g_\tau)\text{-valid for } \Phi(s, \pi)] = \mathbb{P}[\sigma \text{ is } g_\sigma\text{-valid and } \tau \text{ is } g_\tau\text{-valid for } (s, \pi)].$$

Proof. Let Φ be a formula such that (σ, τ) is (g_σ, g_τ) -valid for Φ . Then the total number of pairs (s, π) with $\Phi = \Phi(s, \pi)$ such that σ is g_σ -valid and τ is g_τ -valid for (s, π) equals

$$\prod_{x \in V} \prod_{c, c' \in \{\text{red}, \text{blue}\}} |(\{x\} \times [d_x]) \cap g_\sigma^{-1}(c) \cap g_\tau^{-1}(c')|!,$$

a term that is independent of Φ . □

A *profile* \mathcal{C} consists of two maps $g_\sigma, g_\tau : B \rightarrow \{\text{blue}, \text{red}\}$ and a set $\Gamma \subset g_\sigma^{-1}(\text{blue}) \cap g_\tau^{-1}(\text{red})$ such that $|g_\sigma^{-1}(\text{red})| = |g_\tau^{-1}(\text{red})| = \lambda n$ and such that $\text{red}_x(g_\sigma), \text{red}_x(g_\tau) \leq 3k$ for all $x \in V$.

Let \mathcal{C} be a profile. Moreover, let $\tau \in \{0, 1\}^V$, let s be a signature, and let π be a configuration. We say that (σ, τ, s, π) is *\mathcal{C} -valid* if the following conditions are satisfied.

1. σ, τ are g_σ, g_τ -valid for (s, π) .
2. Let $(x, l) \in g_\sigma^{-1}(\text{blue}) \cap g_\tau^{-1}(\text{red})$. Let $(i, j) = \pi(x, l)$. Then $(x, l) \in \Gamma$ iff $\Phi(s, \pi)_i$ is σ -critical.

In words, this means that (σ, τ, s, π) is \mathcal{C} -valid if σ, τ are solutions of the formula $\Phi(s, \pi)$ under which the colors assigned to the literals by g_σ, g_τ “work out” (i.e., a ball is red iff π puts it in a place such that it supports the clause it occurs in), and if a ball (x, j) belongs to Γ if it supports a clause under τ that is supported by another ball under σ .

Let \mathcal{P} be the set of all profiles. For any $\mathcal{C} \in \mathcal{P}$ and any t let

$$\mu_{\mathcal{C}}(t) = \mathbb{E} \left[\left| \left\{ \tau \in \{0, 1\}^V : \text{dist}(\sigma, \tau) = t \text{ and } (\sigma, \tau, s, \pi) \text{ is } \mathcal{C}\text{-valid} \right\} \right| \right],$$

where the expectation is taken over s, π .

Fact D.5. We have

$$\mu(t) = \frac{\sum_{\mathcal{C} \in \mathcal{P}} \mu_{\mathcal{C}}(t)}{2^{-n} \mathbb{E} \left[\mathcal{Z}'_\beta \right]}. \quad (\text{D.1})$$

Proof. The denominator equals the probability that σ is a NAE-solution that satisfies the first two conditions in Definition 1. Furthermore, $\mu_{\mathcal{C}}(t)$ accounts for the probability that the *pair* (σ, τ) is \mathcal{C} -valid, because for any s, π and any τ there is no more than one profile $\mathcal{C} \in \mathcal{P}$ such that (σ, τ, s, π) is \mathcal{C} -valid. Hence, (D.1) follows from Facts D.3 and D.4. □

We call a profile $\mathcal{C} = (g_\sigma, g_\tau, \Gamma)$ *good* if

$$\frac{1}{n} |g_\sigma^{-1}(\text{red}) \cap g_\tau^{-1}(\text{red})| \in \left[\frac{k}{3 \cdot 2^k}, \frac{3k}{2^k} \right] \quad \text{and} \quad \frac{1}{n} |\Gamma| \in \left[\frac{k^2}{3 \cdot 2^k}, \frac{3k^2}{2^k} \right].$$

Let \mathcal{P}_g be the set of all good profiles, and let $\mathcal{P}_b = \mathcal{P} \setminus \mathcal{P}_g$. Furthermore, let

$$\mu_b(t) = \sum_{\mathcal{C} \in \mathcal{P}_b} \mu_{\mathcal{C}}(t).$$

In Appendix D.3 we are going to show the following.

Proposition D.6. *W.h.p. the degree sequence \mathbf{d} chosen from \mathbf{D} is such that*

$$\sum_{(\frac{1}{2} - 2^{-k/3})n \leq t \leq \frac{n}{2}} \mu_b(t) = o(1).$$

Furthermore, in Appendix D.4 we are going to prove

Proposition D.7. *W.h.p. the degree sequence \mathbf{d} chosen from \mathbf{D} has the following property. Let $\mathcal{C} \in \mathcal{P}_g$ and let $\frac{1}{2} - 2^{-k/3} \leq \alpha \leq \frac{1}{2}$. Then*

$$\mu_{\mathcal{C}}(\alpha n) \leq \exp \left[-c \left(\alpha - \frac{1}{2} \right)^2 n \right] \mu_{\mathcal{C}}(n/2) + \exp(-\Omega(n)).$$

for a certain $c = c(k) > 0$.

We will also need the following fact.

Proposition D.8. *W.h.p. the degree sequence \mathbf{d} chosen from \mathbf{D} is such that $\mu(n/2) \leq \frac{c}{\sqrt{n}} \mathbb{E}[\mathcal{Z}'_\beta]$ for a certain $c = c(k) > 0$.*

Proof. Note that by (D.1) the claim is equivalent to showing

$$\sum_{\mathcal{C} \in \mathcal{P}} \mu_{\mathcal{C}}(n/2) \leq cn^{-1/2} 2^{-n} \mathbb{E}[\mathcal{Z}'_\beta]^2.$$

However, since $\mathbb{E}[\mathcal{Z}'_\beta]$ is the sum of the expectations of indicator random variables over all possible assignments, by expanding $\mathbb{E}[\mathcal{Z}'_\beta]^2$ we arrive at an expression that is a sum over all profiles $\mathcal{C} \in \mathcal{P}$. Then the results follows essentially by performing a term-by-term comparison with the left-hand side of the above inequality. \square

Proposition D.2 is an immediate consequence of (D.1) and Propositions D.6, D.7, and D.8.

D.3 Proof of Proposition D.6

Let Φ be a k -CNF and let $\sigma, \tau \in \{0, 1\}^V$. We say that $(i, j) \in [m] \times [k]$ is σ -red if Φ_{ij} supports Φ_i under σ . Let $\text{red}(\sigma, \Phi)$ be the set of all σ -red pairs (i, j) . We define the term σ -blue and the set $\text{blue}(\sigma, \Phi)$ analogously. Furthermore, let $\Gamma(\sigma, \tau, \Phi)$ be the set of all (i, j) such that $(i, j) \in \text{blue}(\sigma, \Phi) \cap \text{red}(\sigma, \Phi)$ while Φ_i is critical under σ .

Finally, we call the pair $(\sigma, \tau) \in \mathcal{S}(\Phi)^2$ *bad* if $(\frac{1}{2} - 2^{-k/3})n \leq \text{dist}(\sigma, \tau) \leq n/2$ and one of the following conditions holds:

- $|\text{red}(\sigma, \Phi) \cap \text{red}(\tau, \Phi)| \notin \left[\frac{kn}{3 \cdot 2^k}, \frac{3 \cdot kn}{2^k} \right]$, or
- $|\Gamma(\sigma, \tau, \Phi)| \notin \left[\frac{k^2 n}{3 \cdot 2^k}, \frac{3 \cdot k^2 n}{2^k} \right]$.

Lemma D.9. *Let B be the number of bad pairs $(\sigma, \tau) \in \mathcal{S}(\Phi)^2$. Then $\mathbb{E}[B] = \exp(-\Omega(n))$.*

Proof. Let $\sigma = \mathbf{1}$ and let $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$. Let $\mathcal{S}(\alpha)$ be the event that $\sigma, \tau \in \mathcal{S}(\Phi)$. As shown in [3], we have

$$\mathbb{P}[\mathcal{S}] = (1 - 2^{2-k} + 2^{1-k}(\alpha^k + (1-\alpha)^k))^m.$$

Let $R = |\text{red}(\sigma, \Phi) \cap \text{red}(\tau, \Phi)|$. Given that \mathcal{S} occurs, R has a binomial distribution

$$\text{Bin} \left(m, \frac{k(\alpha^k + (1-\alpha)^k)}{(2^{k-1} - 1)(1 - \alpha(1-\alpha)^{k-1} - (1-\alpha)\alpha^{k-1})} \right).$$

For given that σ is a solution, there are a total of $2^k - 2$ ways to choose the signs of the k literals in any clause, and precisely $2k$ ways to choose the signs so that the clause is critical under σ . Given that it is, there are $n^k(1 - \alpha(1-\alpha)^{k-1} - (1-\alpha)\alpha^{k-1})$ ways to choose the actual variables that occur in the clause so as to ensure that τ is a solution, too. (Namely, we have to avoid that either τ and σ differ on the σ -supporting variable only, or that they agree on the σ -supporting variable only; furthermore, the probability that σ, τ differ on a randomly chosen variable is equal to α .) Finally, given that a given clause is σ -critical, the probability that the clause is critical under τ and supported by the same variable as under σ is equal to $\alpha^k + (1-\alpha)^k$ (for σ, τ would either have to agree or disagree on all the k variables).

Further, let $G = |\Gamma(\sigma, \tau, \Phi)|$. Given that \mathcal{S} occurs, G is a binomial variable

$$\text{Bin} \left(m, \frac{k(k-1)(\alpha^2(1-\alpha)^{k-2} + \alpha^{k-2}(1-\alpha)^2)}{(2^{k-1} - 1)(1 - \alpha(1-\alpha)^{k-1} - (1-\alpha)\alpha^{k-1})} \right).$$

For in each σ -critical clause there are $k-1$ ways to choose another literal j to support that clause under τ , and to materialize this choice, τ has to either disagree with σ on the σ -supporting literal and on literal j and agree on all other literals, or the inverse configuration must occur.

It is easily verified that for any $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ we have

$$\mathbb{E}[R|\mathcal{S}] = (1 + o_k(1)) \frac{krn}{2^{2k-2}} \in \left[\frac{kn}{2^k}, \frac{kn}{2^{k-1}} \right],$$

$$\mathbb{E}[G|\mathcal{S}] = (1 + o_k(1)) \frac{k^2 rn}{2^{2k-2}} \in \left[\frac{k^2 n}{2^k}, \frac{k^2 n}{2^{k-1}} \right].$$

As $R, G|\mathcal{S}$ are binomially distributed, Chernoff bounds yield

$$\Pr \left[R \notin \left[\frac{kn}{3 \cdot 2^{k-1}}, \frac{3kn}{2^{k-1}} \right] \right] \leq \exp \left[-\Omega_k \left(\frac{k}{2^k} \right) n \right], \quad (\text{D.2})$$

$$\Pr \left[G \notin \left[\frac{k^2 n}{3 \cdot 2^{k-1}}, \frac{3k^2 n}{2^{k-1}} \right] \right] \leq \exp \left[-\Omega_k \left(\frac{k^2}{2^k} \right) n \right]. \quad (\text{D.3})$$

Since the *total* expected number of pairs of solutions is

$$\mathbb{E}[Z^2] \leq \exp \left[O_k(2^{-k})n \right],$$

the bounds (D.2) and (D.3) imply that $\mathbb{E}[B] \leq \exp(-\Omega(n))$, as claimed. \square

Proposition D.6 is an immediate consequence of Lemma D.9, because the experiment of first choosing \mathbf{d} from the distribution \mathbf{D} and then generating $\Phi_{\mathbf{d}}$ yields precisely the uniform distribution Φ .

D.4 Proof of Proposition D.7

Let $\mathcal{C} = (g_\sigma, g_\tau, \Gamma) \in \mathcal{P}_g$. For $c, c' \in \{\text{red}, \text{blue}\}$ let

$$g_{c,c'} = g_{c,c'}(\mathcal{C}) = |g_\tau^{-1}(c) \cap g_\sigma^{-1}(c')|/n, \text{ and let} \\ \gamma = \gamma(\mathcal{C}) = |\Gamma|/n.$$

Furthermore, for any $\sigma, \tau \in \{0, 1\}^V$ we define

$$\alpha_{c,c'} = \alpha_{c,c'}(\sigma, \tau, \mathcal{C}) = \frac{|\{x \in g_\tau^{-1}(c) \cap g_\sigma^{-1}(c') : \sigma(x) = \tau(x)\}|}{g_{c,c'}n}, \\ \alpha_\Gamma = \alpha_\Gamma(\sigma, \tau, \mathcal{C}) = |\{(x, i) \in \Gamma : \tau(x) = \sigma(x)\}| |\Gamma|, \\ \boldsymbol{\alpha} = \boldsymbol{\alpha}(\sigma, \tau, \mathcal{C}) = (\alpha_{\text{red}, \text{red}}, \alpha_{\text{red}, \text{blue}}, \alpha_{\text{blue}, \text{red}}, \alpha_{\text{blue}, \text{blue}}, \alpha_\Gamma) \in [0, 1]^5.$$

An important observation is that by symmetry, the probability for a pair (σ, τ) to be \mathcal{C} -valid is governed by their “overlap vector” $\boldsymbol{\alpha}$. More precisely, we have

Fact D.10. *Let $\mathcal{C} = (g_\sigma, g_\tau, \Gamma) \in \mathcal{P}_g$. Let $\sigma, \tau, \tau' \in \{0, 1\}^V$ be such that $\boldsymbol{\alpha}(\sigma, \tau, \mathcal{C}) = \boldsymbol{\alpha}(\sigma, \tau', \mathcal{C})$. Then*

$$\mathbb{P}[(\sigma, \tau, \mathbf{s}, \boldsymbol{\pi}) \text{ is } \mathcal{C}\text{-valid}] = \mathbb{P}[(\sigma, \tau', \mathbf{s}, \boldsymbol{\pi}) \text{ is } \mathcal{C}\text{-valid}].$$

Fact D.10 motivates the following definition: for $\boldsymbol{\alpha} = \boldsymbol{\alpha}(\sigma, \tau, \mathcal{C})$ we let

$$p_{\mathcal{C}}(\boldsymbol{\alpha}) = \mathbb{P}[(\sigma, \tau, \mathbf{s}, \boldsymbol{\pi}) \text{ is } \mathcal{C}\text{-valid}].$$

For a real $\alpha \in (0, 1)$ we call a vector $\boldsymbol{\alpha} = (\alpha_{\text{red}, \text{red}}, \dots)$ α -tame if

$$\begin{aligned} |\alpha_{\text{red}, \text{red}} - \alpha| &\leq 10/\sqrt{k}, \\ |\alpha_{\text{red}, \text{blue}} - \alpha| &\leq 2^{-k/3}, \\ |\alpha_{\text{blue}, \text{red}} - \alpha| &\leq 2^{-k/3}, \\ |\alpha_{\text{blue}, \text{blue}} - \alpha| &\leq 2^{-k/2}, \quad \text{and} \\ |\alpha_\Gamma - \alpha| &\leq 100/k. \end{aligned}$$

Let $\mathcal{T}(\alpha)$ be the set of all α -tame vectors. The following lemma shows that we can neglect “overlap vectors” $\boldsymbol{\alpha}$ that are not tame.

Lemma D.11. *Let $\mathcal{C} = (g_\sigma, g_\tau, \Gamma) \in \mathcal{P}_g$. Let W be the number of pairs $(\sigma, \tau) \in \mathcal{S}(\Phi)^2$ with $1 - \alpha = \text{dist}(\sigma, \tau)/n \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ and such that there is a profile \mathcal{C} such that $\boldsymbol{\alpha}(\sigma, \tau, \mathcal{C}) \notin \mathcal{T}(\alpha)$. Then $\mathbb{E}[W] = \exp(-\Omega(n))$.*

The proof of Lemma D.11 is based on a similar first moment argument as in the proof of Lemma D.9. Furthermore, in Section D.5 we will establish the following.

Lemma D.12. *Let $\mathcal{C} = (g_\sigma, g_\tau, \Gamma) \in \mathcal{P}_g$. Let $\boldsymbol{\alpha} \in \mathcal{T}(\alpha)$ for some $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$. Letting $\boldsymbol{\delta} = \boldsymbol{\alpha} - \frac{1}{2}\mathbf{1}$, we have*

$$\begin{aligned} \frac{1}{n} \ln \left(\frac{p_{\mathcal{C}}(\boldsymbol{\alpha})}{p_{\mathcal{C}}(\frac{1}{2}\mathbf{1})} \right) &\leq O_k(k) \cdot [g_{\text{red}, \text{red}}(\delta_{\text{red}, \text{red}}\delta_{\text{blue}, \text{blue}} + \delta_{\text{blue}, \text{blue}}^2) + \gamma(\delta_\Gamma\delta_{\text{blue}, \text{blue}} + \delta_{\text{blue}, \text{blue}}^2)] \\ &\quad + O_k\left(\frac{k^4}{2^k}\right) [\delta_{\text{red}, \text{blue}}\delta_{\text{blue}, \text{blue}} + \delta_{\text{blue}, \text{red}}\delta_{\text{blue}, \text{blue}} + \delta_{\text{blue}, \text{blue}}^2]. \end{aligned}$$

For a number $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ let $p_{\mathcal{C}}(\alpha)$ be the probability that for a random $\tau \in \{0, 1\}^V$ with $\text{dist}(\sigma, \tau) = \alpha n$ we have $\alpha(\sigma, \tau, \mathcal{C}) \in \mathcal{T}(\alpha)$ and (σ, τ, s, π) is \mathcal{C} -valid. We will derive the following consequence of Lemma D.12 in Section D.6.

Corollary D.13. *Suppose that $\alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$ and let \mathcal{C} be a good profile. Then*

$$p_{\mathcal{C}}(\alpha) \leq p_{\mathcal{C}}(1/2) \cdot \exp \left[O_k(k^4/2^k) \cdot \left(\alpha - \frac{1}{2} \right)^2 \cdot n \right] + \exp(-\Omega(n)).$$

Proof of Proposition D.7. By Proposition D.6 and Lemma D.11, for a random \mathbf{d} chosen from \mathbf{D} we have w.h.p.

$$\mu_{\mathcal{C}}(\alpha) \leq \binom{n}{\alpha n} p_{\mathcal{C}}(\alpha) + o(1).$$

Thus, it suffices to estimate $\binom{n}{\alpha n} p_{\mathcal{C}}(\alpha)$. By Stirling's formula and Corollary D.13,

$$\begin{aligned} \frac{1}{n} \ln \left(\frac{\binom{n}{\alpha n} p_{\mathcal{C}}(\alpha)}{\mu_{\mathcal{C}}(1/2)} \right) &\leq \frac{1}{n} \ln \left(\frac{\binom{n}{\alpha n} p_{\mathcal{C}}(\alpha)}{\binom{n}{n/2} p_{\mathcal{C}}(1/2)} \right) \\ &\leq -(4 - o_k(1))(\alpha - 1/2)^2 + \frac{1}{n} \ln \left(\frac{p_{\mathcal{C}}(\alpha)}{p_{\mathcal{C}}(1/2)} \right) \\ &\leq - \left(4 - O_k(\alpha - 1/2) - O_k(k^4/2^k) \right) \cdot \left(\alpha - \frac{1}{2} \right)^2 \\ &= -(4 - o_k(1)) \left(\alpha - \frac{1}{2} \right)^2, \end{aligned}$$

whence the assertion follows for $k \geq k_0$ sufficiently large. \square

D.5 Proof of Lemma D.12

A map $f : [m] \times [k] \rightarrow \{\text{red}, \text{blue}\}$ is called a *coloring* if for each $i \in [m]$ there is at most one $j \in [k]$ such that $f(i, j) = \text{red}$. Let f_{σ}, f_{τ} be colorings. We say that the pair $f = (f_{\sigma}, f_{\tau})$ is *compatible* with a profile $\mathcal{C} = (g_{\sigma}, g_{\tau}, \Gamma)$ if

$$\begin{aligned} |g_{\sigma}^{-1}(c) \cap g_{\tau}^{-1}(c')| &= |f_{\sigma}^{-1}(c) \cap f_{\tau}^{-1}(c')| \quad \text{for any } c, c' \in \{\text{red}, \text{blue}\}, \\ |\Gamma| &= |\{i \in [m] : \exists j \neq l : f_{\sigma}(i, j) = \text{red} \wedge f_{\tau}(i, l) = \text{red}\}|. \end{aligned}$$

Let f be a coloring and let $t : [m] \times [k] \rightarrow \{0, 1\}$ be a map. We call (f, t) *valid* for a signature s if the following two conditions are satisfied:

- for any $i \in [m]$ there exist $j, l \in [k]$ such that $s(i, j)(-1)^{t(i, j)} \neq s(i, l)(-1)^{t(i, l)}$.
- if $f(i, j) = \text{red}$, then for all $l \in [k] \setminus \{j\}$ we have $s(i, j)(-1)^{t(i, j)} \neq s(i, l)(-1)^{t(i, l)}$.

Intuitively, this means that any formula in which the signs are given by s is NAE-satisfied if for all $(i, j) \in [m] \times [k]$ the literal in position (i, j) takes the value $t(i, j)$. Furthermore, for each (i, j) with $f(i, j) = \text{red}$ the literal in position (i, j) supports clause i if the truth values are given by t .

Let $\alpha \in [0, 1]^5$ be a vector. Let $f = (f_{\sigma}, f_{\tau})$ be a pair of colorings. Let $t : [m] \times [k] \rightarrow \{0, 1\}$. We call (f, t) *compatible* with α if

$$\begin{aligned} \alpha_{c, c'} &= \frac{|t^{-1}(1) \cap f_{\sigma}^{-1}(c) \cap f_{\tau}^{-1}(c')|}{|f_{\sigma}^{-1}(c) \cap f_{\tau}^{-1}(c')|} \quad \text{for all } c, c' \in \{\text{red}, \text{blue}\}, \text{ and} \\ \alpha_{\Gamma} &= \frac{|t^{-1}(1) \cap \{(i, l) \in [m] \times [k] : \exists j \neq l : f_{\sigma}(i, j) = \text{red} \wedge f_{\tau}(i, l) = \text{red}\}|}{|\{i \in [m] : \exists j \neq l : f_{\sigma}(i, j) = \text{red} \wedge f_{\tau}(i, l) = \text{red}\}|}. \end{aligned}$$

Let $t : [m] \times [k] \rightarrow \{0, 1\}$ be uniformly distributed, and let

$$q_f(\alpha) = \mathbb{P}_{s,t} [(f, t) \text{ is valid for } s \mid (f, t) \text{ is compatible with } \alpha].$$

Fact D.14. Suppose that f is compatible with a profile \mathcal{C} . Then for any α we have $p_{\mathcal{C}}(\alpha) = q_f(\alpha)$.

Proof. Let $t : [m] \times [k]$ be such that (f, t) is compatible with α . Let $\tau \in \{0, 1\}^V$ be such that $\alpha = \alpha(\sigma, \tau, \mathcal{C})$. Let Π be the set of all $\pi : B \rightarrow [m] \times [k]$ such that $t(\pi(x, i)) = \tau(x)$ for all $x \in V, i \in [d_x]$. Then Π consists of all π that map the right “type” of “ball” to each position (i, j) . Therefore,

$$\begin{aligned} |\Pi| = & ((\alpha_{\Gamma} g_{\Gamma} n)! (1 - \alpha_{\Gamma} g_{\Gamma} n)! \cdot ((\alpha_{\text{red,red}} g_{\text{red,red}} n)! (1 - \alpha_{\text{red,red}} g_{\text{red,red}} n)! \\ & \cdot ((\alpha_{\text{red,blue}} g_{\text{red,blue}} n - \alpha_{\Gamma} g_{\Gamma} n)! ((1 - \alpha_{\text{red,blue}}) g_{\text{red,blue}} n - (1 - \alpha_{\Gamma}) g_{\Gamma} n)! \\ & \cdot ((\alpha_{\text{blue,red}} g_{\text{blue,red}} n - \alpha_{\Gamma} g_{\Gamma} n)! ((1 - \alpha_{\text{blue,red}}) g_{\text{blue,red}} n - (1 - \alpha_{\Gamma}) g_{\Gamma} n)! \\ & \cdot ((\alpha_{\text{blue,blue}} g_{\text{blue,blue}} n - \alpha_{\Gamma} g_{\Gamma} n)! ((1 - \alpha_{\text{blue,blue}}) g_{\text{blue,blue}} n - (1 - \alpha_{\Gamma}) g_{\Gamma} n)! . \end{aligned}$$

Hence, $|\Pi|$ is independent of the actual map t , which implies the assertion. \square

Thus, we are left to compute $q_f(\alpha)$ for a fixed pair $f = (f_{\sigma}, f_{\tau})$ of colorings that is compatible with the good profile \mathcal{C} . To facilitate this computation, we simplify the random experiment further. Namely, let

$$\mathcal{R} = \{(i, j) \in [m] \times [k] : f(i, j) \neq (\text{blue}, \text{blue})\}, \quad \mathcal{B} = [m] \times [k] \setminus \mathcal{R}.$$

For maps $t_{\text{red}} : \mathcal{R} \rightarrow \{0, 1\}$ and $t_{\text{blue}} : \mathcal{B} \rightarrow \{0, 1\}$ we let $t_{\text{red}} \cup t_{\text{blue}} : [m] \times [k]$ be the map defined by

$$(i, j) \mapsto \begin{cases} t_{\text{red}}(i, j) & \text{if } (i, j) \in \mathcal{R}, \\ t_{\text{blue}}(i, j) & \text{if } (i, j) \in \mathcal{B}. \end{cases}$$

Furthermore, we say that (f, t_{red}) is *compatible with α* if there exists t_{blue} such that $(f, t_{\text{red}} \cup t_{\text{blue}})$ is compatible with α .

Suppose that (f, t_{red}) is compatible with α . Let $t_{\text{blue}} : \mathcal{B} \rightarrow \{0, 1\}$ be obtained by setting $t_{\text{blue}}(i, j) = 1$ with probability $\alpha_{\text{blue,blue}}$ and $t_{\text{blue}}(i, j) = 0$ with probability $1 - \alpha_{\text{blue,blue}}$ independently for all $(i, j) \in \mathcal{B}$. Furthermore, let

$$q_f(\alpha, t_{\text{red}}) = \mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s \mid (f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is compatible with } \alpha].$$

Fact D.15. Suppose that (f, t_{red}) is compatible with α . Then $q_f(\alpha) = q_f(\alpha, t_{\text{red}})$.

Lemma D.16. Suppose that (f, t_{red}) is compatible with α . There is a number $C = C(k) > 0$ such that

$$q_f(\alpha, t_{\text{red}}) \leq C \cdot \mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s].$$

Proof. We have

$$\begin{aligned} q_f(\alpha, t_{\text{red}}) &= \mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s \mid (f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is compatible with } \alpha] \\ &= \mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s \mid |t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}|] \\ &= \frac{\mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s \wedge |t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}|]}{\mathbb{P} [|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}|]} \\ &= \mathbb{P} [(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s] \cdot \\ &\quad \frac{\mathbb{P} [|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}| \mid (f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s]}{\mathbb{P} [|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}|]}. \end{aligned} \tag{D.4}$$

We claim that

$$\mathbb{P} \left[|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}| \mid (f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s \right] = O(n^{-1/2}). \quad (\text{D.5})$$

For given that $(f, t_{\text{red}} \cup t_{\text{blue}})$ is valid for s , $|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}|$ is the sum of m independent contributions, as the $t_{\text{blue}}(i, j)$ are independent Bernoulli variables for all $(i, j) \in \mathcal{B}$. Furthermore, given $(f, t_{\text{red}} \cup t_{\text{blue}})$ is valid for s for all i such that $\text{red} \notin f_{\sigma}(i \times [k]) \cup f_{\tau}(i \times [k])$ the random variable $\sum_{j \in [k]} t_{\text{blue}}(i, j)$ takes any value between 1 and k with non-zero probability. Therefore, the conditional random variable $|t_{\text{blue}}^{-1}(1)|$ has a local limit theorem, see Lemma A.1, and (D.5) follows.

As the *unconditional* distribution of $|t_{\text{blue}}^{-1}(1)|$ is just a binomial distribution with mean $\alpha_{\text{blue,blue}} |\mathcal{B}|$, we have

$$\mathbb{P} \left[|t_{\text{blue}}^{-1}(1)| = \alpha_{\text{blue,blue}} |\mathcal{B}| \right] = \Omega(n^{-1/2}).$$

Combining this with (D.4) and (D.5) yields the assertion. \square

Combining Facts D.14 and D.15 with Lemma D.16, we obtain

Corollary D.17. *Suppose that (f, t_{red}) is compatible with α . Then*

$$p_C(\alpha) \leq C \cdot \mathbb{P}[(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s].$$

The crucial feature of the term

$$\mathbb{P}[(f, t_{\text{red}} \cup t_{\text{blue}}) \text{ is valid for } s]$$

is that in the underlying random experiment, the clauses are *independent* objects, although there are different “types” of clauses. This independence property allows us to derive the following estimate.

Proposition D.18. *Suppose that (f, t_{red}) is compatible with α . Let \mathcal{V} be the event that $(f, t_{\text{red}} \cup t_{\text{blue}})$ is valid for s . Let $a = \alpha_{\text{blue,blue}}$. Then*

$$\frac{1}{n} \ln \mathbb{P}[\mathcal{V}] = \psi_{\sigma} + \psi_{\Gamma} + \sum_{c, c' \in \{\text{red, blue}\}} \psi_{c, c'}, \quad (\text{D.6})$$

with the ψ s as shown in Figure 1.

Proof. The first summand ψ_{σ} accounts for the probability that $\sigma = 1$ is a NAE-solution and that precisely the clauses i such that $f(i, j) = \text{red}$ for some $j \in [k]$ are 1-critical. There are precisely λn such clauses, and for each of them the probability of being critical with supporting literal (i, j) equals 2^{1-k} . Furthermore, for the $(r - \lambda)n$ other clauses the probability of being non-critical but NAE-satisfied equals $1 - (k + 1)2^{1-k}$. Since these events depend on the signs of the literals only, they occur independently for all clauses, which explains ψ_{σ} .

The $\psi_{\text{red,red}}$ term is derived quite easily as well. The number of positions (i, j) such that $f_{\sigma}(i, j) = f_{\tau}(i, j) = \text{red}$ equals $g_{\text{red,red}}n$. There are precisely $\alpha_{\text{red,red}}g_{\text{red,red}}n$ among these such that $t_{\text{red}}(i, j) = 1$. Each such position (i, j) supports its clause under t iff $t(i, l) = 1$ for all $l \in [k] \setminus \{j\}$. By the construction of t , the probability of this event is a^{k-1} . Similarly, the “success probability” is $(1 - a)^{k-1}$ for all (i, j) with $t_{\text{red}}(i, j) = 0$.

The next factor ψ_{Γ} accounts for the number of $(i, j) \in f_{\tau}^{-1}(\text{red}) \cap f_{\sigma}^{-1}(\text{blue})$ such that clause i is σ -critical but supported by another literal $l \neq j$ under σ . Each such clause contains precisely $k - 2$ literals $h \in [k] \setminus \{j, l\}$ such that $f_{\tau}(i, h) = f_{\sigma}(i, h) = \text{blue}$. If $t_{\text{red}}(i, j) = 1$, then $t(i, h) = 0$ for all h , which occurs with probability $(1 - a)^{k-2}$. Similarly, if $t_{\text{red}}(i, j) = 0$, then $t(i, h) = 1$ for all h , the probability of which equals a^{k-2} .

$$\begin{aligned}
\psi_\sigma &= (1-k)\lambda \ln 2 + (r-\lambda) \ln(1-(k+1)2^{1-k}), \\
\psi_{\text{red},\text{red}} &= g_{\text{red},\text{red}}(k-1) [\alpha_{\text{red},\text{red}} \ln(a) + (1-\alpha_{\text{red},\text{red}}) \ln(1-a)], \\
\psi_\Gamma &= \gamma(k-2) [\alpha_\gamma \ln(1-a) + (1-\alpha_\gamma) \ln a], \\
\xi &= g_{\text{red},\text{blue}} - \gamma, \\
\alpha_\xi &= \frac{g_{\text{red},\text{blue}}(\alpha_{\text{red},\text{blue}} - \alpha_\Gamma \gamma)}{\xi}, \\
\psi_{\text{red},\text{blue}} &= -\xi \ln(2^{k-1} - k - 1) + \alpha_\xi \xi \ln \left(1 - a^{k-1} - (1-a)^{k-1} - (k-1)a(1-a)^{k-2} \right) \\
&\quad + (1-\alpha_\xi) \xi \ln \left(1 - a^{k-1} - (1-a)^{k-1} - (k-1)a^{k-2}(1-a) \right), \\
\zeta &= g_{\text{blue},\text{red}} - \gamma, \\
\alpha_\zeta &= \frac{g_{\text{blue},\text{red}}(\alpha_{\text{blue},\text{red}} - \alpha_\Gamma \gamma)}{\zeta}, \\
\psi_{\text{blue},\text{red}} &= \alpha_\zeta \zeta \ln \left(1 - a^{k-1} - (1-a)^{k-1} - (k-1)a(1-a)^{k-2} \right) \\
&\quad + (1-\alpha_\zeta) \zeta \ln \left(1 - a^{k-1} - (1-a)^{k-1} - (k-1)a^{k-2}(1-a) \right), \\
\psi_{\text{blue},\text{blue}} &= (r-2\lambda + g_{\text{red},\text{red}}) \ln \left[1 - \frac{1+k-\eta(a)}{2^k-1} - k-1 \right], \quad \text{where} \\
\eta(a) &= a^k + (1-a)^k + ka(1-a)^{k-1} + ka^{k-1}(1-a) + \\
&\quad k(a(1-a)^{k-1} + (1-a)a^{k-1} + a^k + (1-a)^k + \\
&\quad (k-1)a^{k-2}(1-a)^2 + (k-1)a^2(1-a)^{k-2}).
\end{aligned}$$

Fig. 1. The explicit expressions for Proposition D.18.

The term $\psi_{\text{red},\text{blue}}$ deals with clauses i such that $(i, j) \in f_\tau^{-1}(\text{red}) \cap f_\sigma^{-1}(\text{blue}) \setminus \Gamma$ for some j . The total number of such clauses is ξn . For each of these ξn indices i we have $f_\sigma(i, l) = \text{blue}$ for all $l \in [k]$ (because $(i, j) \notin \Gamma$). Suppose that $t_{\text{red}}(i, j) = 1$. Since clause i is non-critical under $\sigma = \mathbf{1}$, it contains a total of $h \geq 2$ literals whose signs agree with that of literal j . In order for clause i to be supported by literal j under \mathbf{t} , the $h-1$ other literals l whose signs agree with that of literal j must take the value $\mathbf{t}(i, l) = 0$, while the $k-h$ remaining literals l must take value $\mathbf{t}(i, l) = 1$. Summing over h and taking into account the distribution of the signs, we obtain the overall probability in the case $t_{\text{red}}(i, j) = 1$:

$$\sum_{j=2}^{k-2} \frac{2^{\binom{k-1}{j-1}}}{2^k - 2^k - 2} (1-a)^{j-1} a^{k-j} = 1 - a^{k-1} - (1-a)^{k-1} - (k-1)a(1-a)^{k-2}.$$

The case $t_{\text{red}}(i, j) = 0$ is analogous to the above, and a similar argument yields $\psi_{\text{blue},\text{red}}$.

Finally, $\psi_{\text{blue},\text{blue}}$ accounts for all clauses i such that $f_\sigma(i, j) = f_\tau(i, j) = \text{blue}$ for all $j \in [k]$. There are precisely $(r-2\lambda + g_{\text{red},\text{red}})n$ such clauses. Each of them is supposed to be assigned such that under both $\sigma = \mathbf{1}$ and \mathbf{t} at least two literals evaluate to “true” and at least two evaluate to “false”. Given the distribution of the signature \mathbf{s} and of \mathbf{t} , the probability of this event equals $\eta(a)$. However, we are already conditioning on the event that each clause contains at least one literal of either sign (this probability is accounted for by ψ_σ). Hence, the conditional probability of the desired outcome equals $\frac{\eta(a)}{1-(k+1)2^{1-k}}$. Since the clauses are independent, the overall probability is given by $\psi_{\text{blue},\text{blue}}$. \square

Proof of Lemma D.12. The assertion simply follows from Proposition D.18 by Taylor expanding the right hand side of (D.6) around $\frac{1}{2}\mathbf{1}$. \square

D.6 Proof of Corollary D.13

We begin with the following observation, which hinges upon the assumption that we work with a good profile.

Proposition D.19. *There is an absolute constant $c > 0$ such that for a random \mathbf{d} chosen from \mathcal{D} the following is true w.h.p. Let \mathcal{C} be a good profile, let $(\frac{1}{2} - 2^{-k/3}) \leq \alpha \leq \frac{1}{2}$, and let τ be chosen uniformly at random from all assignments such that $\text{dist}(\sigma, \tau) = \alpha n$. Then for any $\delta > 0$ we have*

$$\begin{aligned} \mathbb{P}[|\alpha_{\text{blue}, \text{blue}} - \alpha| > \delta] &\leq \exp(-c\delta^2 n), \\ \mathbb{P}[|\alpha_{\text{red}, \text{blue}} - \alpha| > \delta] &\leq \exp(-c\delta^2 n), \\ \mathbb{P}[|\alpha_{\text{blue}, \text{red}} - \alpha| > \delta] &\leq \exp(-c\delta^2 n), \\ \mathbb{P}[|\alpha_{\text{red}, \text{red}} - \alpha| > \delta] &\leq \exp(-g_{\text{red}, \text{red}}\delta^2 n/k^2), \\ \mathbb{P}[|\alpha_{\Gamma} - \alpha| > \delta] &\leq \exp(-\gamma\delta^2 n/k^2). \end{aligned}$$

Proof. Recall that $\sigma = 1$. By standard monotonicity arguments, we may assume that τ is obtained by letting $\tau(x) = 0$ with probability α and $\tau(x) = 1$ with probability $1 - \alpha$ for all $x \in V$ independently. Furthermore, since by standard arguments the degrees d_x are asymptotically independently Poisson, w.h.p. the degree sequence \mathbf{d} is such that

$$\sum_{x \in V} d_x^2 \leq 10 \left(\frac{1}{n} \sum_{x \in V} d_x \right)^2 n \leq 10(kr)^2 n. \quad (\text{D.7})$$

Hence, we are going to assume that (D.7) is satisfied.

We begin by analyzing $\alpha_{\text{blue}, \text{blue}}$. Switching the value $\tau(x)$ of a single variable $x \in V$ can only alter the random variable $\alpha_{\text{blue}, \text{blue}}$ by $d_x/(g_{\text{blue}, \text{blue}}n)$. Therefore, by Azuma's inequality and (D.7), for any $t > 0$

$$\mathbb{P}[|\alpha_{\text{blue}, \text{blue}} - \mathbb{E}[\alpha_{\text{blue}, \text{blue}}]| > t/(g_{\text{blue}, \text{blue}}n)] \leq \exp \left[-\frac{t^2}{\sum_{x \in V} d_x^2} \right] \leq \exp \left[-\frac{t^2}{10n(kr)^2} \right]. \quad (\text{D.8})$$

Since $g_{\text{blue}, \text{blue}} \leq \frac{1}{2}krn$ for any good profile, (D.8) yields the first inequality.

With respect to $\alpha_{\text{red}, \text{blue}}$, recall that in a good profile each $x \in V$ satisfies $\text{red}_\tau(x) \leq k$ (recall that red_τ depends on the profile \mathcal{C} only). Therefore, Azuma's inequality yields

$$\mathbb{P}[|\alpha_{\text{red}, \text{blue}} - \mathbb{E}[\alpha_{\text{red}, \text{blue}}]| > t/(g_{\text{red}, \text{blue}}n)] \leq \exp \left[-\frac{t^2}{k^2 n} \right]. \quad (\text{D.9})$$

Since $g_{\text{red}, \text{blue}} \geq ckn$ for a certain constant $c > 0$, the second claim follows from (D.9). A similar argument yields the third inequality.

Regarding $\alpha_{\text{red}, \text{red}}$, we recall that given \mathcal{C} we know how many “red/red balls” each variable has. Since \mathcal{C} is good, their total number is $g_{\text{red}, \text{red}}n \leq k^2 2^{-k}n$. In particular, there are no more than $g_{\text{red}, \text{red}}n$ variables that have a “red/red ball” in the first place. Furthermore, switching $\tau(x)$ for a single variable x can alter $\alpha_{\text{red}, \text{red}}$ by at most $k/(g_{\text{red}, \text{red}}n)$, because $\text{red}_\tau(x), \text{red}_\sigma(x) \leq k$ for all x as \mathcal{C} is good. Therefore, by Azuma's inequality

$$\mathbb{P}[|\alpha_{\text{red}, \text{red}} - \mathbb{E}[\alpha_{\text{red}, \text{red}}]| > t/(g_{\text{red}, \text{red}}n)] \leq \exp \left[-\frac{t^2}{k^2 g_{\text{red}, \text{red}}n} \right]. \quad (\text{D.10})$$

(The $g_{\text{red}, \text{red}}$ in the denominator mirrors the fact that no more than $g_{\text{red}, \text{red}}n$ variables have a “red/red ball”.) Setting $t = \delta g_{\text{red}, \text{red}}n$ yields the fourth inequality. The last inequality follows from a similar argument. \square

Finally, Corollary D.13 follows by comparing the bounds on the deviations of the individual components of α from Proposition D.19 with Lemma D.12 and Lemma D.11. \square